**Google.Cloud-Digital-Leader.vJun-2024.by.Ana.243q**

Number: Cloud-Digital-Leader
Passing Score: 800
Time Limit: 120
File Version: 7.0

**Certification: Cloud Digital Leader**
**Certification Full Name: Cloud Digital Leader**

**Exam A**

**QUESTION 1**
You are currently managing workloads running on Windows Server for which your company owns the licenses. Your workloads are only needed during working hours, which allows you to shut down the instances during the weekend. Your
Windows Server licenses are up for renewal in a month, and you want to optimize your license cost.
What should you do?

A. Renew your licenses for an additional period of 3 years. Renew your licenses for an additional period of 3 years. Negotiate a cost reduction with your current hosting provider wherein infrastructure cost is reduced when workloads are not in use
B. Renew your licenses for an additional period of 2 years. Negotiate a cost reduction by committing to an automatic renewal of the licenses at the end of the 2 year period
C. Migrate the workloads to Compute Engine with a bring-your-own-license (BYOL) model
D. Migrate the workloads to Compute Engine with a pay-as-you-go (PAYG) model

**Correct Answer: D**
**Section:**
**Explanation:**
The PAYG model is more convenient because you only pay for usage. And the case describes that the workloads are only run on certain days.

**QUESTION 2**
Your organization runs a distributed application in the Compute Engine virtual machines. Your organization needs redundancy, but it also needs extremely fast communication (less than 10 milliseconds) between the parts of the application in different virtual machines.
Where should your organization locate this virtual machines?

A. In a single zone within a single region
B. In different zones within a single region
C. In multiple regions, using one zone per region
D. In multiple regions, using multiple zones per region

**Correct Answer: B**
**Section:**
**Explanation:**
Multi zone is also redundant within the region and it provides the lowest latency.
Reference link:-
https://cloud.google.com/solutions/best-practices-compute-engine-region-selection

**QUESTION 3**
You decide to migrate your on-premises environment to the cloud. You need to determine which resource components still need to be assigned ownership.
Which two functions are owned by a public cloud provider? (Choose two.)

A. Hardware maintenance
B. Infrastructure architecture
C. Infrastructure deployment automation
D. Hardware capacity management
E. Fixing application security issues

**Correct Answer: A, D**
**Section:**
**Explanation:**
In a shared responsible model, hardware maintence and capacity management cloud provider is the responsible part.

**QUESTION 4**
You are a program manager within a Software as a Service (SaaS) company that offers rendering software for animation studios. Your team needs the ability to allow scenes to be scheduled at will and to be interrupted at any time to restart later. Any individual scene rendering takes less than 12 hours to complete, and there is no service-level agreement (SLA) for the completion time for all scenes. Results will be stored in a global Cloud Storage bucket. The compute resources are not bound to any single geographical location. This software needs to run on Google Cloud in a cost-optimized way.
What should you do?

A. Deploy the application on Compute Engine using preemptible instances
B. Develop the application so it can run in an unmanaged instance group
C. Create a reservation for the minimum number of Compute Engine instances you will use
D. Start more instances with fewer virtual centralized processing units (vCPUs) instead of fewer instances with more vCPUs

**Correct Answer: A**
**Section:**
**Explanation:**

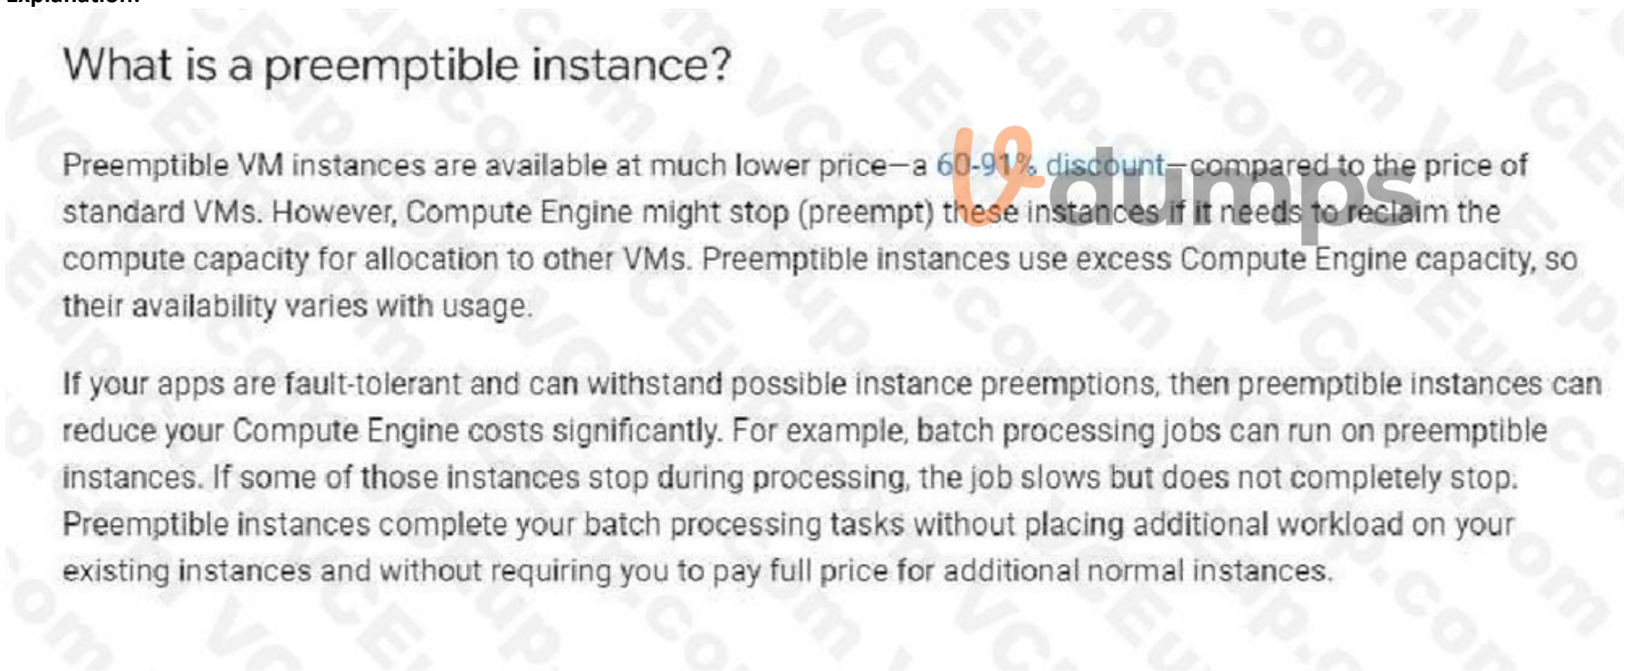


https://cloud.google.com/compute/docs/instances/preemptible

**QUESTION 5**
Your manager wants to restrict communication of all virtual machines with internet access; with resources in another network; or with a resource outside Compute Engine. It is expected that different teams will create new folders and projects in the near future.
How would you restrict all virtual machines from having an external IP address?

A. Define an organization policy at the root organization node to restrict virtual machine instances from having an external IP address
B. Define an organization policy on all existing folders to define a constraint to restrict virtual machine instances from having an external IP address

C. Define an organization policy on all existing projects to restrict virtual machine instances from having an external IP address
D. Communicate with the different teams and agree that each time a virtual machine is created, it must be configured without an external IP address

**Correct Answer: A**
**Section:**
**Explanation:**
Reference: https://cloud.google.com/resource-manager/docs/organization-policy/overview

**QUESTION 6**
Your multinational organization has servers running mission-critical workloads on its premises around the world. You want to be able to manage these workloads consistently and centrally, and you want to stop managing infrastructure.
What should your organization do?

A. Migrate the workloads to a public cloud
B. Migrate the workloads to a central office building
C. Migrate the workloads to multiple local co-location facilities
D. Migrate the workloads to multiple local private clouds

**Correct Answer: A**
**Section:**
**Explanation:**
Only public cloud offers to centrally manage the infra. for Pvt cloud it may not be possible toget same Pvt Cloud provider across the globe.

**QUESTION 7**
Your organization stores highly sensitive data on-premises that cannot be sent over the public internet. The data must be processed both on-premises and in the cloud.
What should your organization do?

A. Configure Identity-Aware Proxy (IAP) in your Google Cloud VPC network
B. Create a Cloud VPN tunnel between Google Cloud and your data center
C. Order a Partner Interconnect connection with your network provider
D. Enable Private Google Access in your Google Cloud VPC network

**Correct Answer: C**
**Section:**
**Explanation:**
After the service provider provisions your connection, you can start passing traffic between your networks by using the service provider's network.
Reference: https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partneroverview

**QUESTION 8**
Your company's development team is building an application that will be deployed on Cloud Run.
You are designing a CI/CD pipeline so that any new version of the application can be deployed in the fewest number of steps possible using the CI/CD pipeline you are designing. You need to select a storage location for the images of the application after the CI part of your pipeline has built them.
What should you do?

A. Create a Compute Engine image containing the application
B. Store the images in Container Registry
C. Store the images in Cloud Storage
D. Create a Compute Engine disk containing the application

**Correct Answer: B**
**Section:**
**Explanation:**
Reference: https://cloud.google.com/container-registry/docs/pushing-and-pulling

**QUESTION 9**
Each of the three cloud service models - infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS) - offers benefits between flexibility and levels of management by the cloud provider and the customer.
Why would SaaS be the right choice of service model?

A. You want a balance between flexibility for the customer and the level of management by the cloud provider
B. You want to minimize the level of management by the customer
C. You want to maximize flexibility for the customer.
D. You want to be able to shift your emphasis between flexibility and management by the cloud provider as business needs change

**Correct Answer: B**
**Section:**
**Explanation:**
Benefits of SaaS
The main benefit of SaaS is that it offloads all infrastructure and application management to the SaaS vendor.
Reference: https://www.ibm.com/cloud/learn/iaas-paas-saas

**QUESTION 10**
As your organization increases its release velocity, the VM-based application upgrades take a long time to perform rolling updates due to OS boot times. You need to make the application deployments faster.
What should your organization do?

A. Migrate your VMs to the cloud, and add more resources to them
B. Convert your applications into containers
C. Increase the resources of your VMs
D. Automate your upgrade rollouts

**Correct Answer: B**
**Section:**
**Explanation:**

**QUESTION 11**
Your organization uses Active Directory to authenticate users. Users' Google account access must be removed when their Active Directory account is terminated.
How should your organization meet this requirement?

A. Configure two-factor authentication in the Google domain
B. Remove the Google account from all IAM policies
C. Configure BeyondCorp and Identity-Aware Proxy in the Google domain
D. Configure single sign-on in the Google domain

**Correct Answer: D**
**Section:**
**Explanation:**
Configure single sign-on in the Google domain
Single sign-on: Whenever a user needs to authenticate, Google Cloud delegates the authentication to Active Directory by using the Security Assertion Markup Language (SAML) protocol. This delegation ensures that only Active Directory manages user credentials and that any applicable policies or multi-factor authentication (MFA) mechanisms are being enforced. For a sign-on to succeed.


https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-introductionReference Link- https://cloud.google.com/architecture/identity/single-sign-on

**QUESTION 12**

Which Google Cloud product gives you a consistent platform for multi-cloud application deployments and extends other Google Cloud services to your environment?

A. Google Kubernetes Engine
B. Virtual Public Cloud
C. Compute Engine
D. Anthos

**Correct Answer: D**
**Section:**
**Explanation:**

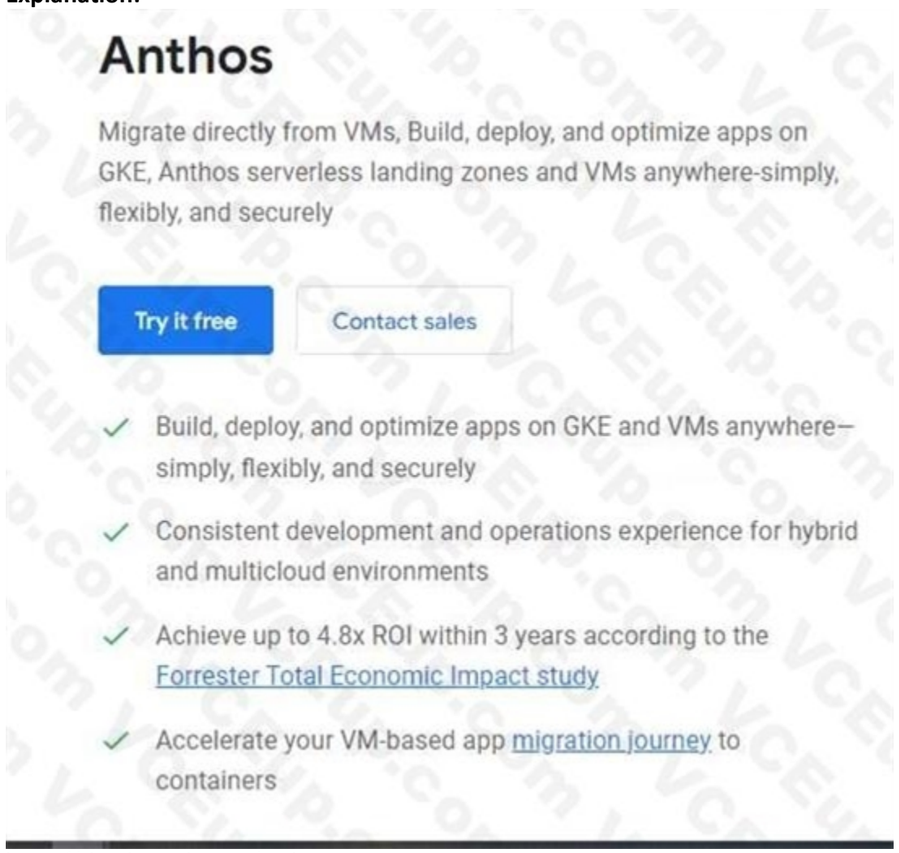


https://cloud.google.com/anthos

**QUESTION 13**
Your organization is developing an application that will manage payments and online bank accounts located around the world. The most critical requirement for your database is that each transaction is handled consistently.
Your organization anticipates almost unlimited growth in the amount of data stored.
Which Google Cloud product should your organization choose?

A. Cloud SQL
B. Cloud Storage
C. Firestore
D. Cloud Spanner

**Correct Answer: D**
**Section:**
**Explanation:**
Features of Cloud Spanner
Reference: https://k21academy.com/google-cloud/cloud-sql-vs-cloud-spanner/

Google Cloud SQL is a fully managed service offered by Google Cloud Platform. **Google Cloud SQL is a MySQL database inside Google Cloud**. There is no need to install, maintain and create admin accounts because it is fully managed by Google Cloud. It helps you create, modify, configure and utilize a relational database, same as MySQL.

Google sends constant updates and adds new features to its services to fulfil the business requirements of its users.

Let's have a look at the improvements done in Google Cloud SQL.

- Google deliberately increased the storage space to **100 Gigabytes** which was 10 Gigabytes earlier.
- The modified version is loaded with the capacity of **16GB RAM** to run instances hassle-free.
- Increased RAM has helped users to keep **four times more cache** than earlier.
- Now Google provides both **replicated and non-replicated databases**.

**QUESTION 14**
Your organization wants an economical solution to store data such as files, graphical images, and videos and to access and share them securely.
Which Google Cloud product or service should your organization use?

A. Cloud Storage
B. Cloud SQL
C. Cloud Spanner
D. BigQuery

**Correct Answer: A**
**Section:**
**Explanation:**
- Google Storage is GCP's version of AWS Simple Storage Service (S3) and an S3 bucket would be equivalent to a Google Storage bucket across the two clouds

Despite many external solutions for digital files, some people still store their photos, videos, and content files on their desktop or laptop. The only problem with this method is that your computer can quickly become cluttered with thousands of files. It slows your prized piece of hardware (computer) down.

When you want to find a digital file you probably *expect* that file to come flying up on your screen in an instant. Yet -- anyone who keeps a lot of photos on a computer knows it can take minutes, sometimes hours, to find one - even if you keep it on your desktop. It's just not all that convenient to store things this way. Most importantly, just storing these digital files on a desktop leaves them vulnerable to viruses, damage, or theft. Folks who rely on this also generally don't have a back-up plan.

**QUESTION 15**

Your organization wants to predict the behavior of visitors to its public website. To do that, you have decided to build a machine learning model. Your team has database-related skills but only basic machine learning skills, and would like to use those database skills.
Which Google Cloud product or feature should your organization choose?

A. BigQuery ML
B. LookML
C. TensorFlow
D. Cloud SQL

**Correct Answer: A**
**Section:**
**Explanation:**
Reference: https://cloud.google.com/architecture/predicting-customer-propensity-to-buy

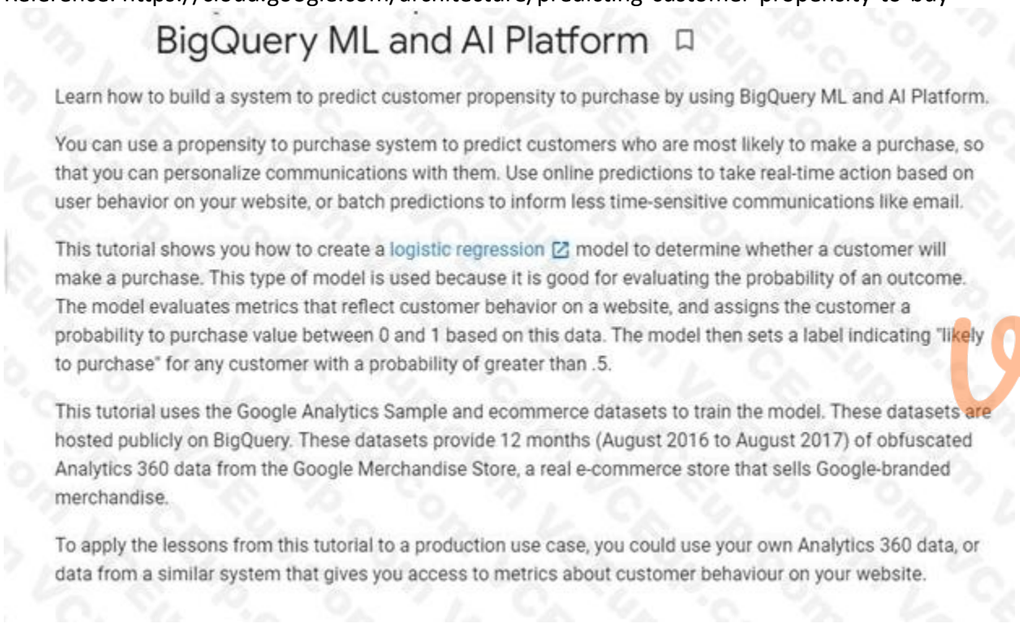

BigQuery ML and AI Platform

Learn how to build a system to predict customer propensity to purchase by using BigQuery ML and AI Platform.

You can use a propensity to purchase system to predict customers who are most likely to make a purchase, so that you can personalize communications with them. Use online predictions to take real-time action based on user behavior on your website, or batch predictions to inform less time-sensitive communications like email.

This tutorial shows you how to create a logistic regression model to determine whether a customer will make a purchase. This type of model is used because it is good for evaluating the probability of an outcome. The model evaluates metrics that reflect customer behavior on a website, and assigns the customer a probability to purchase value between 0 and 1 based on this data. The model then sets a label indicating "likely to purchase" for any customer with a probability of greater than .5.

This tutorial uses the Google Analytics Sample and ecommerce datasets to train the model. These datasets are hosted publicly on BigQuery. These datasets provide 12 months (August 2016 to August 2017) of obfuscated Analytics 360 data from the Google Merchandise Store, a real e-commerce store that sells Google-branded merchandise.

To apply the lessons from this tutorial to a production use case, you could use your own Analytics 360 data, or data from a similar system that gives you access to metrics about customer behaviour on your website.

**QUESTION 16**
Your organization is developing an application that will capture a large amount of data from millions of different sensor devices spread all around the world. Your organization needs a database that is suitable for worldwide, high-speed data storage of a large amount of unstructured data.
Which Google Cloud product should your organization choose?

A. Firestore
B. Cloud Data Fusion
C. Cloud SQL
D. Cloud Bigtable

**Correct Answer: D**
**Section:**
**Explanation:**
Reference: https://cloud.google.com/bigtable
Cloud Bigtable is a sparsely populated table that can scale to billions of rows and thousands of columns, enabling you to store terabytes or even petabytes of data. A single value in each row is indexed; this value is known as the row key.

Bigtable is ideal for storing very large amounts of singlekeyed data with very low latency. It supports high read and write throughput at low latency, and it is an ideal data source for MapReduce operations.
Bigtable is exposed to applications through multiple client libraries, including a supported extension to the Apache HBase library for Java. As a result, it integrates with the existing Apache ecosystem of open-source Big Data software.
Bigtable's powerful back-end servers offer several key advantages over a self-managed HBase installation:
Incredible scalability. Bigtable scales in direct proportion to the number of machines in your cluster. A self-managed HBase installation has a design bottleneck that limits the performance after a certain threshold is reached. Bigtable does not have this bottleneck, so you can scale your cluster up to handle more reads and writes.
Simple administration. Bigtable handles upgrades and restarts transparently, and it automatically maintains high data durability. To replicate your data, simply add a second cluster to your instance, and replication starts automatically. No more managing replicas or regions; just design your table schemas, and Bigtable will handle the rest for you.
Cluster resizing without downtime. You can increase the size of a Bigtable cluster for a few hours to handle a large load, then reduce the cluster's size againóall without any downtime. After you change a cluster's size, it typically takes just a few minutes under load for Bigtable to balance performance across all of the nodes in your cluster.

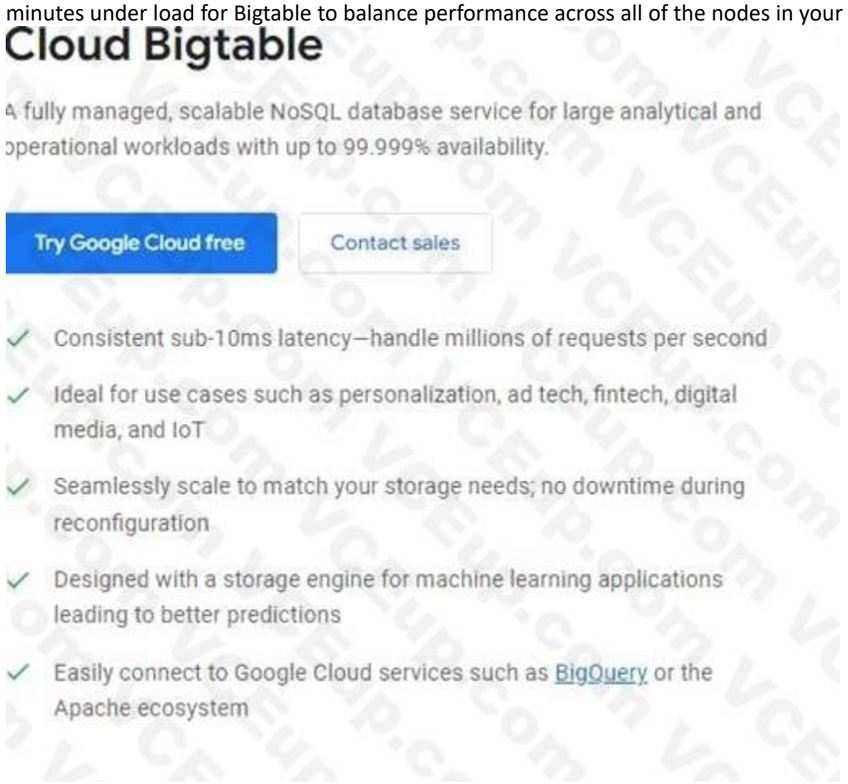


**QUESTION 17**
Your organization needs to build streaming data pipelines. You don't want to manage the individual servers that do the data processing in the pipelines. Instead, you want a managed service that will automatically scale with the amount of data to be processed.
Which Google Cloud product or feature should your organization choose?

A. Pub/Sub
B. Dataflow
C. Data Catalog
D. Dataprep by Trifacta

**Correct Answer: B**
**Section:**
**Explanation:**
Reference: https://cloud.google.com/dataflow/docs/guides/deploying-a-pipeline

## Deploying a pipeline

Send feedback

★ This document explains in detail how Dataflow deploys and runs a pipeline, and covers advanced topics like optimization and load balancing. If you are looking for a step-by-step guide on how to create and deploy your first pipeline, use Dataflow's quickstarts for Java, Python, Go, or templates.

After you construct and test your Apache Beam pipeline, you can use the Dataflow managed service to deploy and execute it. Once on the Dataflow service, your pipeline code becomes a Dataflow job.

The Dataflow service fully manages Google Cloud services such as Compute Engine and Cloud Storage to run your Dataflow job, automatically spinning up and tearing down the necessary resources. The Dataflow service provides visibility into your job through tools like the Dataflow monitoring interface and the Dataflow command-line interface.

★ You can control some aspects of how the Dataflow service runs your job by setting execution parameters in your pipeline code. For example, the execution parameters specify whether the steps of your pipeline run on worker virtual machines, on the Dataflow service backend, or locally.

In addition to managing Google Cloud resources, the Dataflow service automatically performs and optimizes many aspects of distributed parallel processing. These include the following:

- **Parallelization and distribution.** Dataflow automatically partitions your data and distributes your worker code to Compute Engine instances for parallel processing. For more information, see parallelization and distribution.
- **Optimization.** Dataflow uses your pipeline code to create an execution graph that represents your pipeline's `PCollection`s and transforms, and optimizes the graph for the most efficient performance and resource usage. Dataflow also automatically optimizes potentially costly operations, such as data aggregations. For more information, see Fusion optimization and Combine optimization.

Reference link- https://cloud.google.com/dataflow/docs/guides/deploying-a-pipeline



**QUESTION 18**
Your organization is building an application running in Google Cloud. Currently, software builds, tests, and regular deployments are done manually, but you want to reduce work for the team. Your organization wants to use Google Cloud managed solutions to automate your build, testing, and deployment process.
Which Google Cloud product or feature should your organization use?

A. Cloud Scheduler
B. Cloud Code
C. Cloud Build
D. Cloud Deployment Manager

**Correct Answer: C**
**Section:**
**Explanation:**
Deploy your application to App Engine using the gcloud app deploy command. This command automatically builds a container image by using the Cloud Build service and then deploys that image to the App Engine flexible environment.
Reference: https://cloud.google.com/appengine/docs/flexible/nodejs/testing-and-deploying-yourapp

**QUESTION 19**
Which Google Cloud product can report on and maintain compliance on your entire Google Cloud organization to cover multiple projects?

A. Cloud Logging
B. Identity and Access Management

C. Google Cloud Armor
D. Security Command Center

**Correct Answer: D**
**Section:**
**Explanation:**
Security Command Center is a centralized security and risk management platform for yourGoogle Cloud resources. It is a single tool that offers a variety of security features including:
1. Gain centralized visibility and control
2. Discover misconfigurations and vulnerabilities
3. Report on and maintain compliance
4. Detect threats targeting your Google Cloud assets
https://cloud.google.com/security-command-center

**QUESTION 20**
Your organization needs to establish private network connectivity between its on-premises network and its workloads running in Google Cloud. You need to be able to set up the connection as soon as possible.
Which Google Cloud product or feature should you use?

A. Cloud Interconnect
B. Direct Peering
C. Cloud VPN
D. Cloud CDN

**Correct Answer: A**
**Section:**
**Explanation:**
Private Google Access for on-premises hosts provides a way for on-premises systems to connect to Google APIs and services by routing traffic through a Cloud VPN tunnel.
Reference: https://cloud.google.com/vpc/docs/configure-private-google-access-hybrid

**QUESTION 21**
Your organization is developing a mobile app and wants to select a fully featured cloud-based compute platform for it.
Which Google Cloud product or feature should your organization use?

A. Google Kubernetes Engine
B. Firebase
C. Cloud Functions
D. App Engine

**Correct Answer: B**
**Section:**
**Explanation:**
Reference: https://cloud.google.com/appengine
Firebase is Google's mobile development platform that empowers you to quickly build and grow your app

**QUESTION 22**
Your team has developed a machine learning model for your customer. The test results indicate very strong predictive capability. The model is then deployed in production. Evaluation of the predictions in production show that they are off by a pronounced margin. What is the issue and how can you solve for it?

A. The model is under fitted. Train with less data.
B. The model is over fitted. Add more features to the model to fix it.

C. The model is fine since the test results are good. Fix the production of incoming data.
D. The model is overfitted. Train with more data.

**Correct Answer: D**
**Section:**
**Explanation:**
If our ML model does well on the training set than on the production set, then we're likely over fitting. Training with more data would be one solution.

**QUESTION 23**
Your large and frequently changing organization's user information is stored in an on-premises LDAP database. The database includes user passwords and group and organization membership.
How should your organization provision Google accounts and groups to access Google Cloud resources?

A. Replicate the LDAP infrastructure on Compute Engine
B. Use the Firebase Authentication REST API to create users
C. Use Google Cloud Directory Sync to create users
D. Use the Identity Platform REST API to create users

**Correct Answer: C**
**Section:**
**Explanation:**
You can run a single instance of Google Cloud Directory Sync to synchronize user accounts and groups to Google Cloud.
Reference: https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-introduction

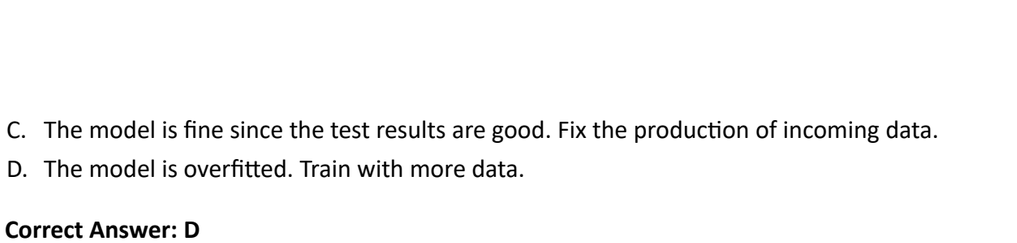


https://support.google.com/a/answer/106368?hl=en

**QUESTION 24**
Your organization recently migrated its compute workloads to Google Cloud. You want these workloads in Google Cloud to privately and securely access your large volume of on-premises data, and you also want to minimize latency.
What should your organization do?

A. Use Storage Transfer Service to securely make your data available to Google Cloud
B. Create a VPC between your on-premises data center and your Google resources
C. Peer your on-premises data center to Google's Edge Network
D. Use Transfer Appliance to securely make your data available to Google Cloud

**Correct Answer: C**
**Section:**
**Explanation:**

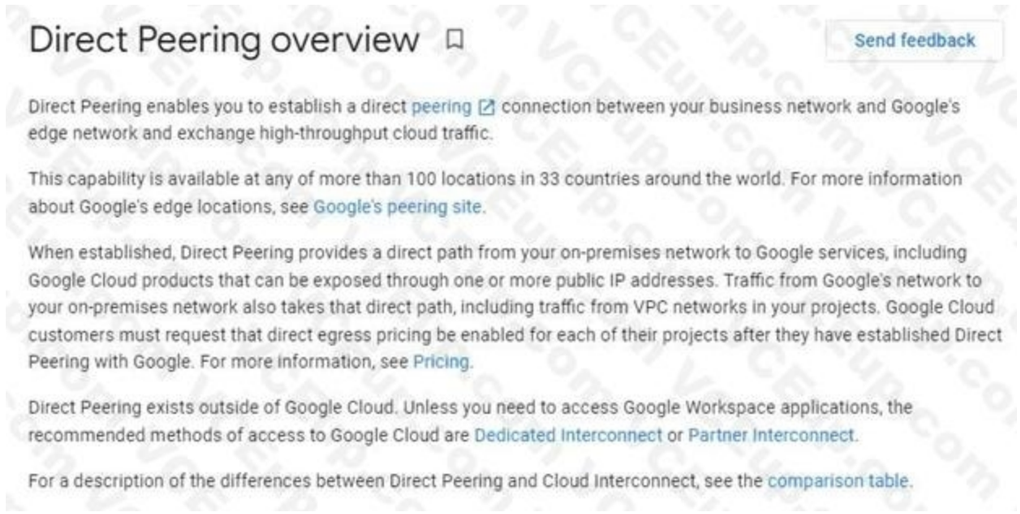

https://cloud.google.com/network-connectivity/docs/direct-peering

**QUESTION 25**
Your organization consists of many teams. Each team has many Google Cloud projects. Your organization wants to simplify the management of identity and access policies for these projects.
How can you group these projects to meet this goal?

A. Group each team's projects into a separate domain
B. Assign labels based on the virtual machines that are part of each team's projects
C. Use folders to group each team's projects
D. Group each team's projects into a separate organization node

**Correct Answer: C**
**Section:**
**Explanation:**

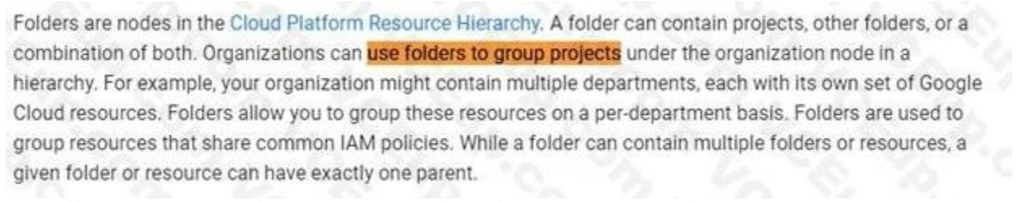

https://cloud.google.com/resource-manager/docs/creating-managing-folders

**QUESTION 26**
Your organization needs to restrict access to a Cloud Storage bucket. Only employees who are based in Canada should be allowed to view the contents.
What is the most effective and efficient way to satisfy this requirement?

A. Deploy the Cloud Storage bucket to a Google Cloud region in Canada
B. Configure Google Cloud Armor to allow access to the bucket only from IP addresses based in Canada
C. Give each employee who is based in Canada access to the bucket
D. Create a group consisting of all Canada-based employees, and give the group access to the bucket

**Correct Answer: D**
**Section:**

**Explanation:**
Reference: https://cloud.google.com/storage/docs/access-control
Because you can use your own private VPN to access the Canada-only bucket from anywhere in the world.

**QUESTION 27**
Your organization is moving an application to Google Cloud. As part of that effort, it needs to migrate the application's working database from another cloud provider to Cloud SQL. The database runs on the MySQL engine. The migration must cause minimal disruption to users. Data must be secured while in transit.
Which should your organization use?

A. BigQuery Data Transfer Service
B. MySQL batch insert
C. Database Migration Service
D. Cloud Composer

**Correct Answer: C**
**Section:**
**Explanation:**
Reference: https://aws.amazon.com/dms/

**QUESTION 28**
Your organization is developing and deploying an application on Google Cloud. Tracking your Google Cloud spending needs to stay as simple as possible.
What should you do to ensure that workloads in the development environment are fully isolated from production workloads?

A. Apply a unique tag to development resources
B. Associate the development resources with their own network
C. Associate the development resources with their own billing account
D. Put the development resources in their own project

**Correct Answer: D**
**Section:**
**Explanation:**
Reference: https://www.deps.co/blog/google-cloud-platform-good-bad-ugly/Project resources are components that are necessary for successful project implementation. Theyinclude people, equipment, money, time, knowledge ñ basically, anything that you may require fromthe project planning to the project delivery phases.

**QUESTION 29**
Your company is running the majority of its workloads in a co-located data center. The workloads are running on virtual machines (VMs) on top of a hypervisor and use either Linux or Windows server editions. As part of your company's transformation strategy, you need to modernize workloads as much as possible by adopting cloud-native technologies. You need to migrate the workloads into Google Cloud.
What should you do?

A. Export the VMs into VMDK format, and import them into Compute Engine
B. Export the VMs into VMDK format, and import them into Google Cloud VMware Engine
C. Migrate the workloads using Migrate for Compute Engine
D. Migrate the workloads using Migrate for Anthos

**Correct Answer: D**
**Section:**
**Explanation:**
Anthos: Anthos lets you build, deploy, and manage applications anywhere in a secure, consistent manner. You can modernize existing applications running on virtual machines while deploying cloudnative apps on containers

in an increasingly hybrid and multi-cloud world.

**QUESTION 30**
Your organization is running all its workloads in a private cloud on top of a hypervisor. Your organization has decided it wants to move to Google Cloud as quickly as possible. Your organization wants minimal changes to the current environment, while using the maximim amount of managed services Google offers.
What should your organization do?

A. Migrate the workloads to Google Cloud VMware Engine
B. Migrate the workloads to Compute Engine
C. Migrate the workloads to Bare Metal Solution
D. Migrate the workloads to Google Kubernetes Engine

**Correct Answer: B**
**Section:**
**Explanation:**
Migrate for Compute Engine enables you to lift and shift workloads at scale to Google Cloud Compute Engine with minimal changes and risk.
Reference: https://dataintegration.info/simplify-vm-migrations-with-migrate-for-compute-engineas-a-service

**QUESTION 31**
Your organization is releasing its first publicly available application in Google Cloud. The application is critical to your business and customers and requires a 2-hour SLA.
How should your organization set up support to minimize costs?

A. Enroll in Premium Support
B. Enroll in Enhanced Support
C. Enroll in Standard Support
D. Enroll in Basic Support

**Correct Answer: B**
**Section:**
**Explanation:**
Reference: https://www.secureauth.com/enhanced-support-offering/
SecureAuth is dedicated to providing the industry-leading enhanced support ensuring the long term success of your SecureAuth SaaS IAM deployment

SecureAuth is dedicated to providing the **industry-leading** enhanced **support** ensuring the long term success of your SecureAuth SaaS IAM deployment.

While our basic support offers industry leading coverage and response times for some customers, SecureAuth protects critical applications meaning delays and extended downtime is simply not an option. For these customers our **Enhanced Support** offerings provide **24*7 coverage** and the most responsive and **complete SLA**'s available. That's why we offer three different levels of support, so you can choose the level of support that best works for your needs.



**QUESTION 32**
Your organization offers public mobile apps and websites. You want to migrate to a Google Cloudbased solution for checking and maintaining your users' usernames and passwords and controlling their access to different resources based on their identity.
Which should your organization choose?

A. VPN tunnels
B. Identity Platform
C. Compute Engine firewall rules
D. Private Google Access

**Correct Answer: B**
**Section:**
**Explanation:**
An identity platform is a modern solution for managing the identities of users and devices in a centralized fashion.
Reference: https://www.okta.com/blog/2021/07/what-is-an-identityplatform/#:~:text=An%20identity%20platform%20is%20a,%2C%20integrations%2C%20and%20platform%20services

An identity platform is a modern solution for managing the identities of users and devices in a centralized fashion. It enables organizations to securely authorize workforce and customer users to access their ecosystem using access management tools, programmable components, integrations, and platform services.

Today's organizations have a wide range of identity requirements. As they expand, embrace new innovations, and meet new customer demands, they need an identity solution that can grow with them. And as they grow, these organizations are also looking for offerings that centralize and consolidate identity, reducing the need for multiple access management, governance, and authentication products that don't necessarily talk to each other.

In short, although the concept of identity platforms is still new, the growth of disruptive technologies and highly personalized products and services has made a platform approach the key for effectively unifying identity management. But how did we get here?

**QUESTION 33**
Which Google Cloud service or feature lets you build machine learning models using Standard SQL and data in a data warehouse?

A. BigQuery ML
B. TensorFlow
C. AutoML Tables
D. Cloud Bigtable ML

**Correct Answer: A**
**Section:**
**Explanation:**
BigQuery ML lets you create and execute machine learning models in BigQuery using standard SQL queries.
Reference: https://cloud.google.com/bigqueryml/docs/introduction#:~:text=BigQuery%20ML%20lets%20you%20create,the%20need%20to%20move%20data

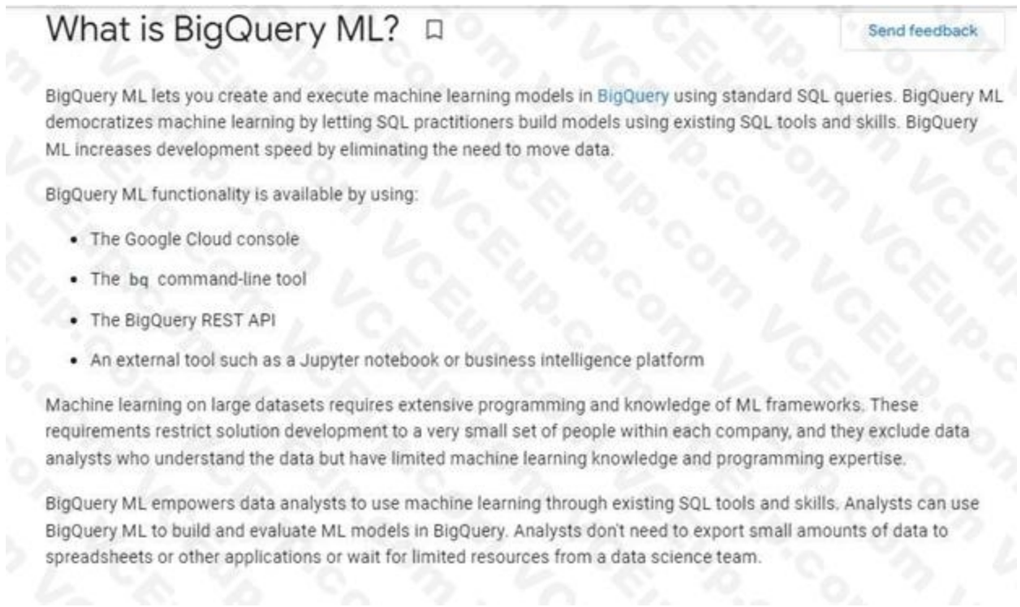

What is BigQuery ML?

Send feedback

BigQuery ML lets you create and execute machine learning models in BigQuery using standard SQL queries. BigQuery ML democratizes machine learning by letting SQL practitioners build models using existing SQL tools and skills. BigQuery ML increases development speed by eliminating the need to move data.

BigQuery ML functionality is available by using:

- The Google Cloud console
- The `bq` command-line tool
- The BigQuery REST API
- An external tool such as a Jupyter notebook or business intelligence platform

Machine learning on large datasets requires extensive programming and knowledge of ML frameworks. These requirements restrict solution development to a very small set of people within each company, and they exclude data analysts who understand the data but have limited machine learning knowledge and programming expertise.

BigQuery ML empowers data analysts to use machine learning through existing SQL tools and skills. Analysts can use BigQuery ML to build and evaluate ML models in BigQuery. Analysts don't need to export small amounts of data to spreadsheets or other applications or wait for limited resources from a data science team.

https://cloud.google.com/bigquery-ml/docs/introduction

**QUESTION 34**
Your organization runs an application on virtual machines in Google Cloud. This application processes incoming images. This activity takes hours to create a result for each image. The workload for this application normally stays at a certain baseline level, but at regular intervals it spikes to a much greater workload. Your organization needs to control the cost to run this application.
What should your organization do?

A. Purchase committed use discounts for the baseline load
B. Purchase committed use discounts for the expected spike load
C. Leverage sustained use discounts for your virtual machines
D. Run the workload on preemptible VM instances

**Correct Answer: C**
**Section:**
**Explanation:**
The idea of the Sustained Use discount is that the longer you run a VM instance in any given month, the bigger discount you will get from the list price.
Reference: https://www.parkmycloud.com/blog/google-sustained-use-discounts/

**QUESTION 35**
Your organization is developing a plan for migrating to Google Cloud.
What is a best practice when initially configuring your Google Cloud environment?

A. Create a project via Google Cloud Console per department in your company
B. Define your resource hierarchy with an organization node on top
C. Create projects based on tem members' requests
D. Make every member of your company the project owner

**Correct Answer: B**
**Section:**
**Explanation:**

The Organization resource is the root node of the Google Cloud resource hierarchy and all resources that belong to an organization are grouped under the organization node. This provides central visibility and control over every resource that belongs to an organization.
Reference link- https://cloud.google.com/resource-manager/docs/cloud-platform-resource-hierarchy

**QUESTION 36**
Your organization runs many workloads in different Google Cloud projects, each linked to the same billing account. Each project's workload costs can vary from month to month, but the overall combined cost of all projects is relatively stable. Your organization needs to optimize its cost.
What should your organization do?

A. Purchase a commitment per project for each project's usual minimum
B. Create a billing account per project, and link each project to a different billing account
C. Turn on committed use discount sharing, and create a commitment for the combined usage
D. Move all workloads from all different projects into one single consolidated project

**Correct Answer: C**
**Section:**
**Explanation:**
Turn on committed use discount sharing, and create a commitment for the combined usage Sharing your committed use discounts across all your projects reduces the overhead of managing discounts on a per-project basis, and maximizes your savings by pooling all your discounts across your projects' resource usage. If you have multiple projects that share the same Cloud Billing account, you can enable committed use discount sharing so all of your projects within that Cloud Billing account share all of your committed use discount contracts. Your sustained use discounts are also pooled at the same time. That is, sustained use discounts are calculated using the total resources across these projects, rather than just the resources within a single project.

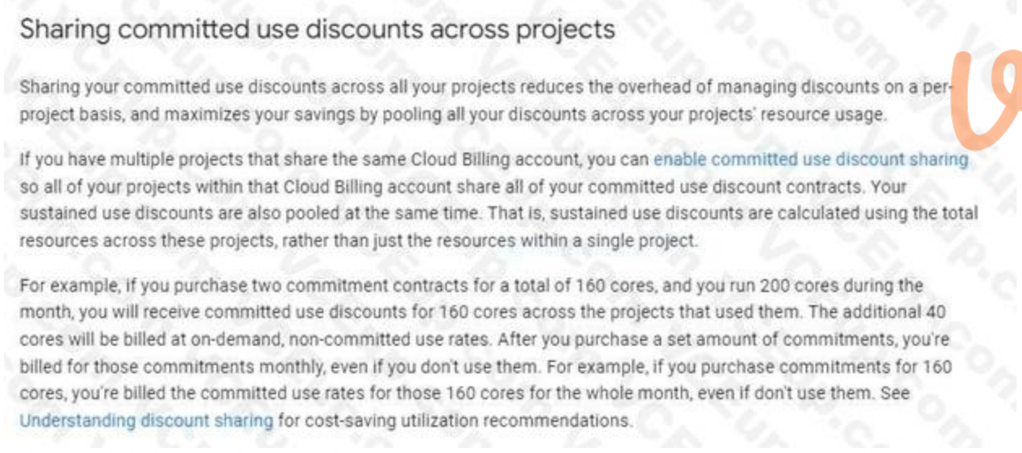
Sharing committed use discounts across projects

Sharing your committed use discounts across all your projects reduces the overhead of managing discounts on a per-project basis, and maximizes your savings by pooling all your discounts across your projects' resource usage.

If you have multiple projects that share the same Cloud Billing account, you can enable committed use discount sharing so all of your projects within that Cloud Billing account share all of your committed use discount contracts. Your sustained use discounts are also pooled at the same time. That is, sustained use discounts are calculated using the total resources across these projects, rather than just the resources within a single project.

For example, if you purchase two commitment contracts for a total of 160 cores, and you run 200 cores during the month, you will receive committed use discounts for 160 cores across the projects that used them. The additional 40 cores will be billed at on-demand, non-committed use rates. After you purchase a set amount of commitments, you're billed for those commitments monthly, even if you don't use them. For example, if you purchase commitments for 160 cores, you're billed the committed use rates for those 160 cores for the whole month, even if don't use them. See Understanding discount sharing for cost-saving utilization recommendations.

Reference link- https://cloud.google.com/compute/docs/instances/signing-up-committed-usediscounts#sharing_committed_use_discounts_across_projects

**QUESTION 37**
How should a multinational organization that is migrating to Google Cloud consider security and privacy regulations to ensure that it is in compliance with global standards?

A. Comply with data security and privacy regulations in each geographical region
B. Comply with regional standards for data security and privacy, because they supersede all international regulations
C. Comply with international standards for data security and privacy, because they supersede all regional regulations
D. Comply with regional data security regulations, because they're more complex than privacy standards

**Correct Answer: A**
**Section:**
**Explanation:**

Comply with data security and privacy regulations in each geographical region For a multi-national corporation, they need to abide not just by international laws, but also regional laws where they do business.

**QUESTION 38**
Your company has recently acquired three growing startups in three different countries. You want to reduce overhead in infrastructure management and keep your costs low without sacrificing security and quality of service to your customers.
How should you meet these requirements?

A. Host all your subsidiaries' services on-premises together with your existing services.
B. Host all your subsidiaries' services together with your existing services on the public cloud.
C. Build a homogenous infrastructure at each subsidiary, and invest in training their engineers.
D. Build a homogenous infrastructure at each subsidiary, and invest in hiring more engineers.

**Correct Answer: B**
**Section:**
**Explanation:**
Host all your subsidiaries' services together with your existing services on the public cloud.

**QUESTION 39**
What is the difference between Standard and Coldline storage?

A. Coldline storage is for data for which a slow transfer rate is acceptable.
B. Standard and Coldline storage have different durability guarantees.
C. Standard and Coldline storage use different APIs.
D. Coldline storage is for infrequently accessed data.

**Correct Answer: D**
**Section:**
**Explanation:**
Reference: https://www.msp360.com/resources/blog/google-cloud-nearline-storage-vs-coldline-vsstandard/Google Cloud Coldline is a new cold-tier storage for archival data with access frequency of less thanonce per year. Unlike other cold storage options, Nearline has no delays prior to data access, so nowit is the leading solution among competitors.

The main characteristics of Coldline are as follows:

- SLA guarantees 99% data availability (of ten thousand hours, data can be offline for up to 100 hours).
- Monthly fee per GB stored is between $0.004-$.0.014, depending on a region.
- The minimum storing period is 90 days. If you delete data earlier, you have to pay for the remaining time. For example, if you had uploaded 100GB of data and then deleted it after 30 days, you need to pay extra $1.4 as an early deletion fee ($0.004 (or $0.014) * 100 GB * 2 months).
- Data retrieval fee is mandatory and costs $0.05 per GB.

Its low price and a long minimum storing period make Nearline the best solution for data that is unlikely to be accessed more than once a year, if ever:

- Data archive.
- Disaster recovery storage.
- Outdated backups storage.

When using Coldline or Nearline Storage, you also pay more for requests, e.g., retrieval of metadata or download commands.

**QUESTION 40**
What would provide near-unlimited availability of computing resources without requiring your organization to procure and provision new equipment?

A. Public cloud
B. Containers
C. Private cloud
D. Microservices

**Correct Answer: A**
**Section:**
**Explanation:**
Reference: https://cloud.google.com/docs/overview

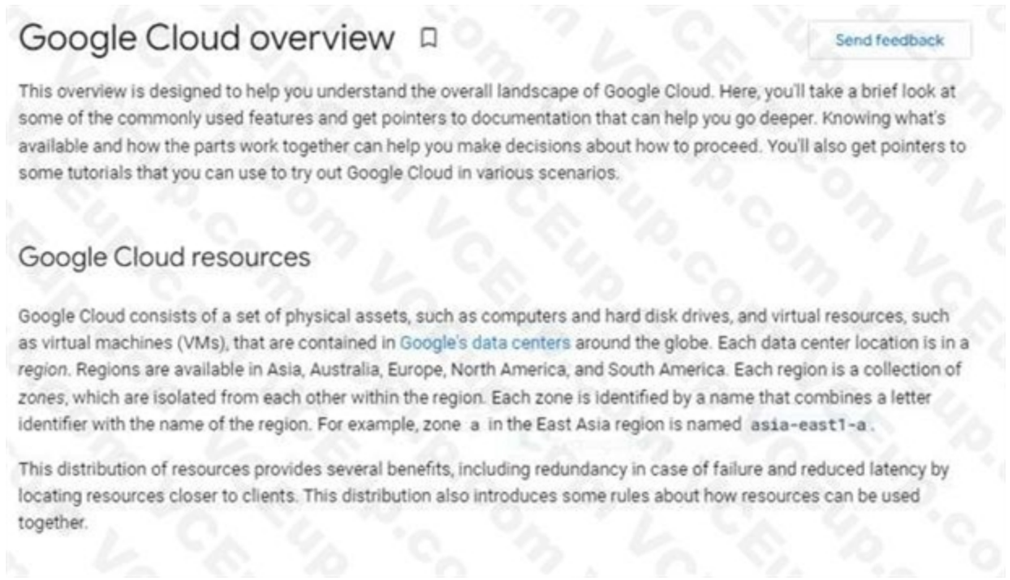


**QUESTION 41**
You are a program manager for a team of developers who are building an event-driven application to allow users to follow one another's activities in the app. Each time a user adds himself as a follower of another user, a write occurs in the real-time database.
The developers will develop a lightweight piece of code that can respond to database writes and generate a notification to let the appropriate users know that they have gained new followers. The code should integrate with other cloud services such as Pub/Sub, Firebase, and Cloud APIs to streamline the orchestration process. The application requires a platform that automatically manages underlying infrastructure and scales to zero when there is no activity.
Which primary compute resource should your developers select, given these requirements?

A. Google Kubernetes Engine
B. Cloud Functions
C. App Engine flexible environment
D. Compute Engine

**Correct Answer: B**
**Section:**
**Explanation:**
Reference: https://firebase.google.com/docs/functions/use-cases

Cloud Functions gives developers access to Firebase and Google Cloud events, along with scalable computing power to run code in response to those events. While it's expected that Firebase apps will use Cloud Functions in unique ways to meet their unique requirements, typical use cases might fall into these areas:

- Notify users when something interesting happens.
- Perform database sanitization and maintenance.
- Execute intensive tasks in the cloud instead of in your app.
- Integrate with third-party services and APIs.

Review the use cases and examples for each category that interests you, and then proceed to our Get Started tutorial or to specific how-to guides for authentication events, analytics events, and more. See the `eventType` API reference for the complete list of supported event types.

## Notify users when something interesting happens

Developers can use Cloud Functions to keep users engaged and up to date with relevant information about an app. Consider, for example, an app that allows users to follow one another's activities in the app. Each time a user adds themselves as a follower of another user, a write occurs in the Realtime Database. Then this write event could trigger a function to create Firebase Cloud Messaging (FCM) notifications to let the appropriate users know that they have gained new followers.

**QUESTION 42**
Your company has been using a shared facility for data storage and will be migrating to Google Cloud. One of the internal applications uses Linux custom images that need to be migrated.
Which Google Cloud product should you use to maintain the custom images?

A. App Engine flexible environment
B. Compute Engine
C. App Engine standard environment
D. Google Kubernetes Engine

**Correct Answer: B**
**Section:**
**Explanation:**
Reference: https://cloud.google.com/compute/docs/images/create-delete-deprecate-privateimagesA custom image is a boot disk image that you own and control access to. Use custom images for the following tasks:
Import a virtual disk to Compute Engine from your on-premises environment or from VMs that are running on your local workstation or on another cloud platform. You can manually import boot disk images to Compute Engine, but one disk at a time.

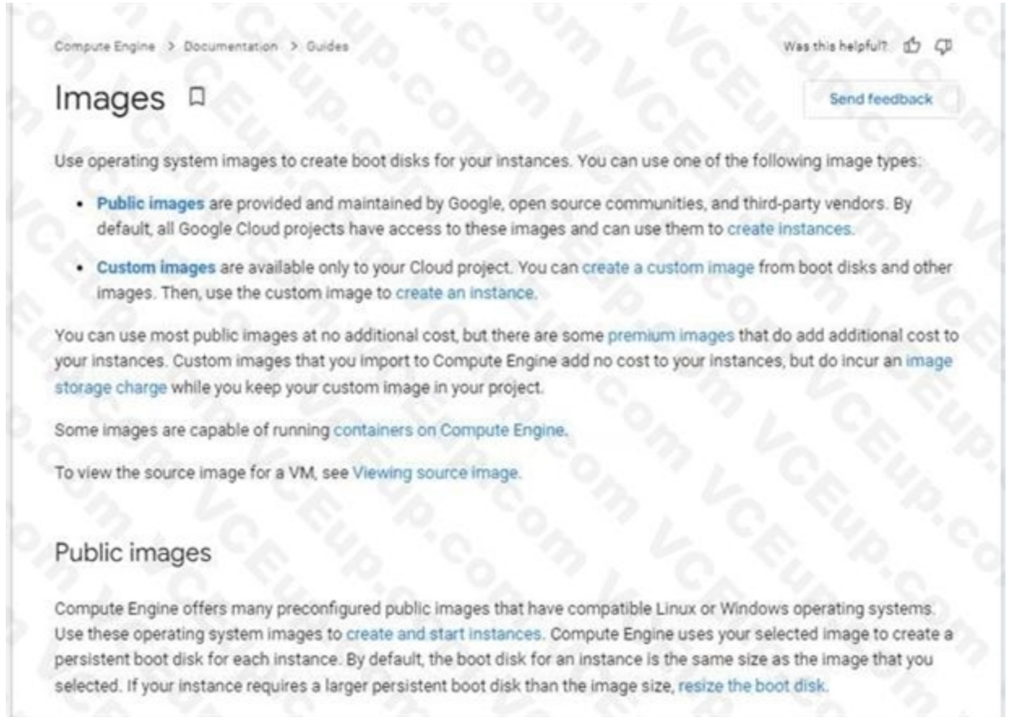
Compute Engine > Documentation > Guides

Was this helpful?

# Images

Send feedback

Use operating system images to create boot disks for your instances. You can use one of the following image types:

- **Public images** are provided and maintained by Google, open source communities, and third-party vendors. By default, all Google Cloud projects have access to these images and can use them to create instances.
- **Custom images** are available only to your Cloud project. You can create a custom image from boot disks and other images. Then, use the custom image to create an instance.

You can use most public images at no additional cost, but there are some premium images that do add additional cost to your instances. Custom images that you import to Compute Engine add no cost to your instances, but do incur an image storage charge while you keep your custom image in your project.

Some images are capable of running containers on Compute Engine.

To view the source image for a VM, see Viewing source image.

## Public images

Compute Engine offers many preconfigured public images that have compatible Linux or Windows operating systems. Use these operating system images to create and start instances. Compute Engine uses your selected image to create a persistent boot disk for each instance. By default, the boot disk for an instance is the same size as the image that you selected. If your instance requires a larger persistent boot disk than the image size, resize the boot disk.

https://cloud.google.com/compute/docs/images

**QUESTION 43**
Your organization wants to migrate its data management solutions to Google Cloud because it needs to dynamically scale up or down and to run transactional SQL queries against historical data at scale. Which Google Cloud product or service should your organization use?

A. BigQuery
B. Cloud Bigtable
C. Pub/Sub
D. Cloud Spanner

**Correct Answer: D**
**Section:**
**Explanation:**
Reference: https://cloud.google.com/terms/services
Cloud Spanner is a fully-managed, mission-critical relational database service. It is designed to provide a scalable online transaction processing (OLTP) database with high availability and strong consistency at global scale

**QUESTION 44**
Your organization needs to categorize objects in a large group of static images using machine learning. Which Google Cloud product or service should your organization use?

A. BigQuery ML
B. AutoML Video Intelligence
C. Cloud Vision API
D. AutoML Tables

**Correct Answer: C**

**Section:**
**Explanation:**
Reference: https://cloud.google.com/vision
Derive insights from your images in the cloud or at the edge with AutoML Vision or use pre-trained Vision API models to detect emotion, understand text, and more.
Vision API offers powerful pre-trained machine learning models through REST and RPC APIs. Assign labels to images and quickly classify them into millions of predefined categories. Detect objects and faces, read printed and handwritten text, and build valuable metadata into your image catalog.

**QUESTION 45**
Your organization runs all its workloads on Compute Engine virtual machine instances. Your organization has a security requirement: the virtual machines are not allowed to access the public internet. The workloads running on those virtual machines need to access BigQuery and Cloud Storage, using their publicly accessible interfaces, without violating the security requirement.
Which Google Cloud product or feature should your organization use?

A. Identity-Aware Proxy
B. Cloud NAT (network address translation)
C. VPC internal load balancers
D. Private Google Access

**Correct Answer: D**
**Section:**
**Explanation:**
VM instances that only have internal IP addresses (no external IP addresses) can use Private Google Access. They can reach the external IP addresses of Google APIs and services. The source IP address of the packet can be the primary internal IP address of the network interface or an address in an alias IP range that is assigned to the interface. If you disable Private Google Access, the VM instances can no longer reach Google APIs and services; they can only send traffic within the VPC network.

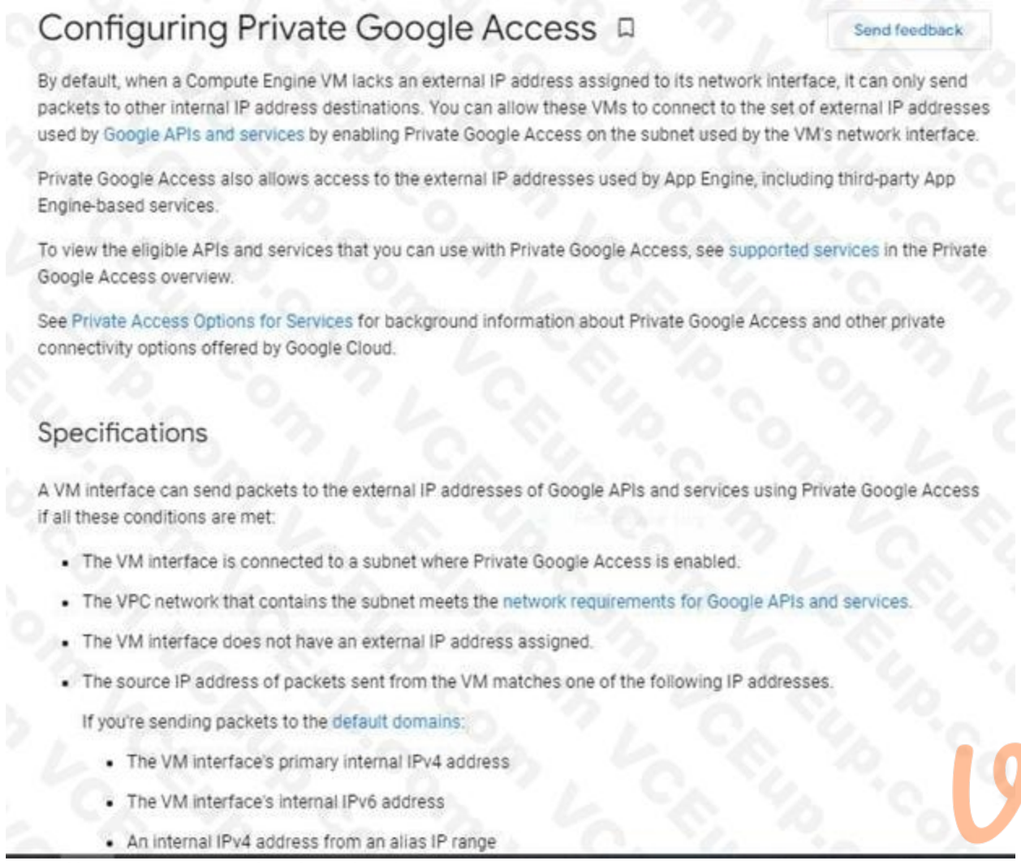
Configuring Private Google Access

Send feedback

By default, when a Compute Engine VM lacks an external IP address assigned to its network interface, it can only send packets to other internal IP address destinations. You can allow these VMs to connect to the set of external IP addresses used by Google APIs and services by enabling Private Google Access on the subnet used by the VM's network interface.

Private Google Access also allows access to the external IP addresses used by App Engine, including third-party App Engine-based services.

To view the eligible APIs and services that you can use with Private Google Access, see supported services in the Private Google Access overview.

See Private Access Options for Services for background information about Private Google Access and other private connectivity options offered by Google Cloud.

Specifications

A VM interface can send packets to the external IP addresses of Google APIs and services using Private Google Access if all these conditions are met:

- The VM interface is connected to a subnet where Private Google Access is enabled.
- The VPC network that contains the subnet meets the network requirements for Google APIs and services.
- The VM interface does not have an external IP address assigned.
- The source IP address of packets sent from the VM matches one of the following IP addresses.

  If you're sending packets to the default domains:

  - The VM interface's primary internal IPv4 address
  - The VM interface's internal IPv6 address
  - An internal IPv4 address from an alias IP range

https://cloud.google.com/vpc/docs/configure-private-google-access

**QUESTION 46**
Which Google Cloud product is designed to reduce the risks of handling personally identifiable information (PII)?

A. Cloud Storage
B. Google Cloud Armor
C. Cloud Data Loss Prevention
D. Secret Manager

**Correct Answer: C**
**Section:**
**Explanation:**
Reference: https://cloud.google.com/blog/products/gcp/take-charge-of-your-sensitive-data-withthe-cloud-dlp-api
Cloud Data Loss Prevention: Fully managed service designed to help you discover, classify, and protect your most sensitive data.

the general availability of the Cloud Data Loss Prevention (DLP) API, a Google Cloud security service that helps you discover, classify and redact sensitive data at rest and in real-time.

When it comes to properly handling sensitive data, the first step is knowing where it exists in your data workloads. This not only helps enterprises more tightly secure their data, it's a fundamental component of reducing risk in today's regulatory environment, where the mismanagement of sensitive information can come with real costs.

The DLP API is a flexible and robust tool that helps identify sensitive data like credit card numbers, social security numbers, names and other forms of personally identifiable information (PII). Once you know where this data lives, the service gives you the option to de-identify that data using techniques like redaction, masking and tokenization. These features help protect sensitive data while allowing you to still use it for important business functions like running analytics and customer support operations. On top of that, the DLP API is designed to plug into virtually any workload—whether in the cloud or on-prem—so that you can easily stream in data and take advantage of our inspection and de-identification capabilities.

**QUESTION 47**
Your organization is migrating to Google Cloud. As part of that effort, it needs to move terabytes of data from on-premises file servers to Cloud Storage. Your organization wants the migration process to be automated and to be managed by Google. Your organization has an existing Dedicated Interconnect connection that it wants to use. Which Google Cloud product or feature should your organization use?

A. Storage Transfer Service
B. Migrate for Anthos
C. BigQuery Data Transfer Service
D. Transfer Appliance

**Correct Answer: A**
**Section:**
**Explanation:**
Reference: https://cloud.google.com/architecture/migration-to-google-cloud-transferring-yourlarge-datasets

Storage Transfer Service for large transfers of on-premises data

Like `gsutil`, Storage Transfer Service for on-premises data enables transfers from network file system (NFS) storage to Cloud Storage. Although `gsutil` can support small transfer sizes (up to 1 TB), Storage Transfer Service for on-premises data is designed for large-scale transfers (up to petabytes of data, billions of files). It supports full copies or incremental copies, and it works on all transfer options listed earlier in Deciding among Google's transfer options. It also has a simple, managed graphical user interface, even non-technically savvy users (after setup) can use it to move data.

Storage Transfer Service for on-premises data is especially useful in the following scenarios:

- You have sufficient available bandwidth to move the data volumes (see the Google Cloud Data Transfer Calculator).
- You support a large base of internal users who might find a command-line tool like `gsutil` challenging to use.
- You need robust error-reporting and a record of all files and objects that are moved.
- You need to limit the impact of transfers on other workloads in your data center (this product can stay under a user-specified bandwidth limit).
- You want to run recurring transfers on a schedule.

| Where you're moving data from | Scenario | Suggested products |
|---|---|---|
| Another cloud provider (for example, Amazon Web Services or Microsoft Azure) to Google Cloud | — | Storage Transfer Service |
| Cloud Storage to Cloud Storage (two different buckets) | — | Storage Transfer Service |
| Your private data center to Google Cloud | Enough bandwidth to meet your project deadline for less than 1 TB of data | `gsutil` |
| Your private data center to Google Cloud | Enough bandwidth to meet your project deadline for more than 1 TB of data | Storage Transfer Service for on-premises data |
| Your private data center to Google Cloud | Not enough bandwidth to meet your project deadline | Transfer Appliance |

https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets

**QUESTION 48**
Your organization needs to analyze data in order to gather insights into its daily operations. You only want to pay for the data you store and the queries you perform. Which Google Cloud product should your organization choose for its data analytics warehouse?

A. Cloud SQL
B. Dataproc
C. Cloud Spanner
D. BigQuery

**Correct Answer: D**
**Section:**
**Explanation:**
BigQuery is an enterprise data warehouse for large amounts of relational structured data Serverless, highly scalable, and cost-effective multicloud data warehouse designed for business agility.

**QUESTION 49**
Your organization wants to run a container-based application on Google Cloud. This application is expected to increase in complexity. You have a security need for fine-grained control of traffic between the containers. You

also have an operational need to exercise fine-grained control over the application's scaling policies.
What Google Cloud product or feature should your organization use?

A. Google Kubernetes Engine cluster
B. App Engine
C. Cloud Run
D. Compute Engine virtual machines

**Correct Answer: A**
**Section:**
**Explanation:**
Google Kubernetes Engine GKE seems a better fit since the requirement is for "security need for finegrained control of traffic between the containers" and "fine-grained control over scaling policies".
Such level of control is easier on GKE that Cloud Run.

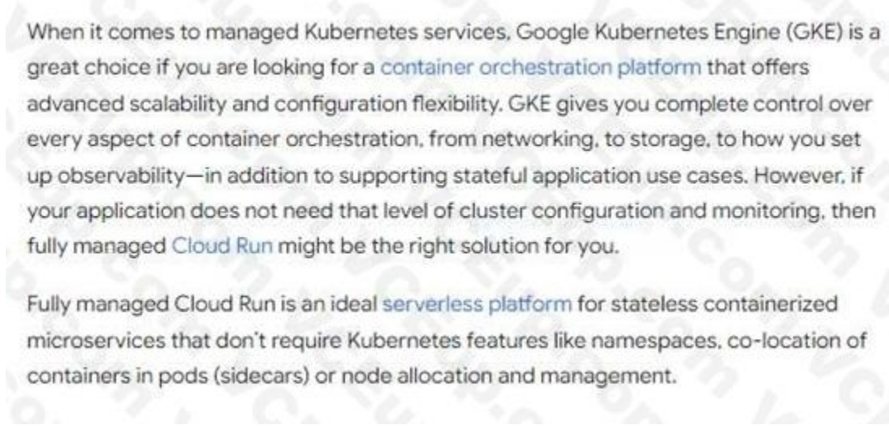
When it comes to managed Kubernetes services, Google Kubernetes Engine (GKE) is a great choice if you are looking for a container orchestration platform that offers advanced scalability and configuration flexibility. GKE gives you complete control over every aspect of container orchestration, from networking, to storage, to how you set up observability—in addition to supporting stateful application use cases. However, if your application does not need that level of cluster configuration and monitoring, then fully managed Cloud Run might be the right solution for you.

Fully managed Cloud Run is an ideal serverless platform for stateless containerized microservices that don't require Kubernetes features like namespaces, co-location of containers in pods (sidecars) or node allocation and management.

Reference link- https://cloud.google.com/blog/products/containers-kubernetes/when-to-usegoogle-kubernetes-engine-vs-cloud-run-for-containers

**QUESTION 50**
Which Google Cloud product or feature makes specific recommendations based on security risks and compliance violations?

A. Google Cloud firewalls
B. Security Command Center
C. Cloud Deployment Manager
D. Google Cloud Armor

**Correct Answer: B**
**Section:**
**Explanation:**
Reference: https://cloud.google.com/security-command-center
Security Command Center is Security and risk management platform for Google Cloud.

Asset discovery and inventory

Discover and view your assets in near-real time across App Engine, BigQuery, Cloud SQL, Cloud Storage, Compute Engine, Cloud Identity and Access Management, Google Kubernetes Engine, and more. Review historical discovery scans to identify new, modified, or deleted assets.

Threat prevention

Understand the security state of your Google Cloud assets. Uncover common web application vulnerabilities such as cross-site scripting or outdated libraries in your web applications running on App Engine, GKE, and Compute Engine. Quickly resolve misconfigurations by clicking directly on the impacted resource and following the prescribed steps on how to fix it.

Threat detection

Detect threats using logs running in Google Cloud at scale. Detect some of the most common container attacks, including suspicious binary, suspicious library, and reverse shell.

Reference link- https://cloud.google.com/security-command-center

**QUESTION 51**
Your organization wants to optimize its use of Google Cloud's discounts on virtual machine-based workloads. You plan to use 200 CPUs constantly for the next 3 years, and you forecast that spikes of up to 300 CPUs will occur approximately 30% of the time. What should you choose?

A. 1-year committed use discount for 200 CPUs
B. 3-year committed use discount for 300 CPUs
C. 3-year committed use discount for 200 CPUs
D. Regular pay-as-you-go pricing

**Correct Answer: C**
**Section:**

**QUESTION 52**
Your organization needs to minimize how much it pays for data traffic from the Google network to the internet. What should your organization do?

A. Choose the Standard network service tier.
B. Choose the Premium network service tier.
C. Deploy Cloud VPN.
D. Deploy Cloud NAT.

**Correct Answer: A**
**Section:**
**Explanation:**

Choose the Standard network service tier. While Premium tier is the default for all egress traffic and offers the highest performance, when cost is a consideration. Standard tier is the more economical.

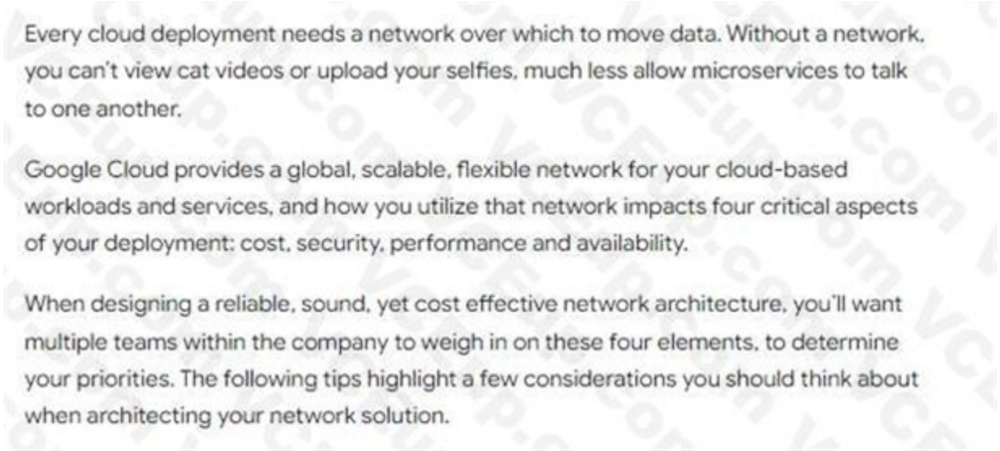

Every cloud deployment needs a network over which to move data. Without a network, you can't view cat videos or upload your selfies, much less allow microservices to talk to one another.

Google Cloud provides a global, scalable, flexible network for your cloud-based workloads and services, and how you utilize that network impacts four critical aspects of your deployment: cost, security, performance and availability.

When designing a reliable, sound, yet cost effective network architecture, you'll want multiple teams within the company to weigh in on these four elements, to determine your priorities. The following tips highlight a few considerations you should think about when architecting your network solution.

https://cloud.google.com/blog/products/networking/networking-cost-optimization-best-practices

**QUESTION 53**
Your organization wants to migrate your on-premises environment to Google Cloud. The on-premises environment consists of containers and virtual machine instances. Which Google Cloud products can help to migrate the container images and the virtual machine disks?

A. Compute Engine and Filestore
B. Artifact Registry and Cloud Storage
C. Dataflow and BigQuery
D. Pub/Sub and Cloud Storage

**Correct Answer: A**
**Section:**
**Explanation:**
Reference: https://cloud.google.com/compute/docs/import/importing-virtual-disks

If you have virtual disks in your on-premises environment with software and configurations that you need (sometimes referred to as *golden disks* or *golden images*), you can save time by importing those virtual disks into Compute Engine and using the resulting image to create virtual machines. The import tool supports most virtual disk file formats, including VMDK and VHD.

If you exported your disk from Compute Engine, you can create images from the disk.

For information about how to create an automated system for migrating several virtual machines (VMs), see Migrating VMs to Compute Engine.

**QUESTION 54**
Your company security team manages access control to production systems using an LDAP directory group.
How is this access control managed in the Google Cloud production project?

A. Assign the proper role to the Service Account in the project's IAM Policy
B. Grant each user the roles/iam.serviceAccountUser role on a service account that exists in the Google Group.
C. Assign the proper role to the Google Group in the project's IAM Policy.
D. Create the project in a folder with the same name as the LDAP directory group.

**Correct Answer: C**
**Section:**
**Explanation:**
Reference: https://cloud.google.com/blog/products/identity-security/achieving-identity-and-accessgovernance-on-google-cloud

When businesses shift from solely on-premises deployments to using cloud-based services, identity management can become more complex. This is especially true when it comes to hybrid and multi-cloud identity management.

Cloud Identity and Access Management (IAM) offers several ways to manage identities and roles in Google Cloud. One particularly important identity management task is identity and access governance (IAG): ensuring that your identity and access permissions are managed effectively, securely, and correctly. A major step in achieving IAG is designing an architecture that suits your business needs and also allows you to satisfy your compliance requirements. To manage the entire enterprise identity lifecycle you must consider the following core tasks:



**QUESTION 55**
Your organization wants to be sure that is expenditures on cloud services are in line with the budget.
Which two Google Cloud cost management features help your organization gain greater visibility into its cloud resource costs? (Choose two.)

A. Billing dashboards
B. Resource labels
C. Sustained use discounts
D. Financial governance policies
E. Payments profile

**Correct Answer: A, B**
**Section:**
**Explanation:**


A label is a key-value pair that helps you organize your Google Cloud resources. You can attach a label to each resource, then filter the resources based on their labels. Information about labels is forwarded to the billing system, so you can break down your billed charges by label.
Reference link- https://cloud.google.com/cost-management

**QUESTION 56**
Your organization needs to process large amounts of data from an online application that operates continuously. You do not want to be required to provision infrastructure or create server clusters.
What should your organization choose?

A. Compute Engine with BigQuery
B. Dataproc
C. Google Kubernetes Engine with Cloud Bigtable
D. Dataflow

**Correct Answer: D**
**Section:**
**Explanation:**
You do not want to be required to provision infrastructure or create server clusters. Because Unified stream and batch data processing that's serverless, fast, and cost-effective.
Reference link- https://cloud.google.com/dataflow

**QUESTION 57**
Your organization needs to ensure that the Google Cloud resources of each of your departments are segregated from one another. Each department has several environments of its own: development, testing, and production.
Which strategy should your organization choose?

A. Create a project per department, and create a folder per environment in each project.
B. Create a folder per department, and create a project per environment in each folder.
C. Create a Cloud Identity domain per department, and create a project per environment in each domain.
D. Create a Cloud Identity domain per environment, and create a project per department in each domain.

**Correct Answer: B**
**Section:**
**Explanation:**

Folders are nodes in the Cloud Platform Resource Hierarchy. A folder can contain projects, other folders, or a combination of both. Organizations can use folders to group projects under the organization node in a hierarchy. For example, your organization might contain multiple departments, each with its own set of Google Cloud resources. Folders allow you to group these resources on a per-department basis. Folders are used to group resources that share common IAM policies. While a folder can contain multiple folders or resources, a given folder or resource can have exactly one parent.

```
# Template for new folder & new project

folder_resource = {
  'name': 'new-folder',
  'type': 'gcp-types/cloudresourcemanager-v2:folders',
  'properties': {
     'parent': 'organizations/99999',
     'displayName': 'new-folder'
  }
}

project_resource = {
  'name': 'new-project',
  'type': 'clouresourcemanager.v1.project',
  'metadata': { 'dependsOn': ['new-folder'] },
  'properties': {
     'name': 'new-project',
     'parent': {
        'type': 'folder',
        # HERE it is -- the problem!
        'id': '$(ref.new-folder.name)'
     }
  }
}

return { 'resources': [folder_resource, project_resource] }
```

Reference link- https://cloud.google.com/resource-manager/docs/creating-managing-folders
Reference link- https://stackoverflow.com/questions/59460623/how-to-create-a-folder-a-project-under-it-with-deployment-manager-google-cloud

**QUESTION 58**
Your organization is defining the resource hierarchy for its new application in Google Cloud. You need separate development and production environments. The production environment will be deployed in Compute Engine in two regions.
Which structure should your organization choose?

A. Create a single project for all environments. Use labels to segregate resources by environment.
B. Create a single project for all environments. Use tags to segregate resources by environment.
C. Create one project for the development environment and one project for the production environment.
D. Create two projects for the development environment and two projects for the production environment (one for each region).

**Correct Answer: C**
**Section:**
**Explanation:**
Many organizations have separate development and production environments so they can build and test new features without disturbing production traffic. In Optimizely, you can create separate projects for each environment to help with governance.
With separate development and production projects, your organization can safely build and QA experiments and Personalization campaigns in a development environment before deploying to production. This approach allows multiple stakeholders in your organization to act as gatekeepers for running new experiments in production.


Reference link- https://support.optimizely.com/hc/en-us/articles/4410284353805-Set-up-projects-for-development-and-production-environments

**QUESTION 59**
Your organization meant to purchase a 3-year Committed Use Discount, but accidentally purchased a 1-year Committed Use Discount instead. What should your organization do?

A. Contact your financial institution.
B. Contact Trust and Safety.
C. Contact Cloud Billing Support.
D. Contact Technical Support.

**Correct Answer: C**
**Section:**
**Explanation:**

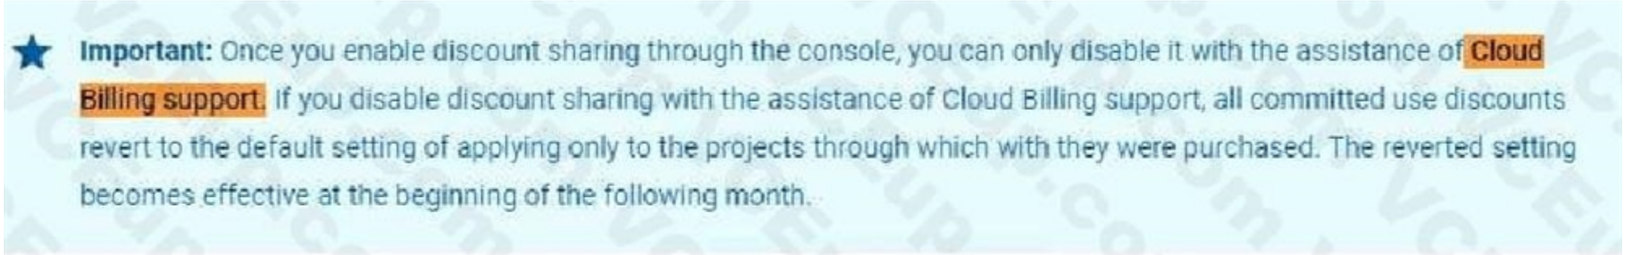


https://cloud.google.com/compute/docs/instances/signing-up-committed-use-discounts

**QUESTION 60**
Your organization needs to allow a production job to have access to a BigQuery dataset. The production job is running on a Compute Engine instance that is part of an instance group.
What should be included in the IAM Policy on the BigQuery dataset?

A. The Compute Engine instance group
B. The project that owns the Compute Engine instance
C. The Compute Engine service account
D. The Compute Engine instance

**Correct Answer: C**
**Section:**
**Explanation:**
When an identity calls a Google Cloud API, BigQuery requires that the identity has the appropriate permissions to use the resource. You can grant permissions by granting roles to a user, a group, or a service account.
Reference link- https://cloud.google.com/bigquery/docs/access-control

**QUESTION 61**
Your team is publishing research results and needs to make large amounts of data available to other researchers within the professional community and the public at minimum cost.
How should you host the data?

A. Use a Cloud Storage bucket and enable "Requester Pays.'
B. Use a Cloud Storage bucket and provide Signed URLs for the data files.
C. Use a Cloud Storage bucket and set up a Cloud Interconnect connection to allow access to the data.
D. Host the data on-premises. and set up a Cloud Interconnect connection to allow access to the data.

**Correct Answer: A**
**Section:**
**Explanation:**
Enabling Requester Pays is useful, for example, if you have a lot of data you want to make available to users, but you don't want to be charged for their access to that data.
Reference link- https://cloud.google.com/storage/docs/requester-pays

**QUESTION 62**
Your company needs to segment Google Cloud resources used by each team from the others. The teams' efforts are changing frequently, and you need to reduce operational risk and maintain cost visibility. Which approach does Google recommend?

A. One project per team.
B. One organization per team.
C. One project that contains all of each team's resources.
D. One top-level folder per team.

**Correct Answer: A**
**Section:**
**Explanation:**
Reference: https://cloud.google.com/security/infrastructure/design
The Teams need to segmented to have visibility on the resources each team consumes

**QUESTION 63**

How do Migrate for Compute Engine and Migrate for Anthos differ?

A. Unlike Migrate for Anthos, Migrate for Compute Engine assumes that the migration source is VMware vSphere.
B. Migrate for Compute Engine charges for ingress, but Migrate for Anthos does not.
C. Migrate for Compute Engine is closed source, and Migrate for Anthos is open source.
D. Migrate for Anthos migrates to containers, and Migrate for Compute Engine migrates to virtual machines.

**Correct Answer: D**
**Section:**
**Explanation:**
Reference: https://cloud.google.com/migrate/anthos
Migrate workloads to Compute Engine with Migrate for Compute Engine. Migrate from Compute Engine to containers with Migrate for Anthos and GKE.
This method makes sense, for instance, in cases where you want to conduct a data-center migration and migrate all workloads into Compute Engine, and only at a second stage selectively modernize suitable workloads to containers.

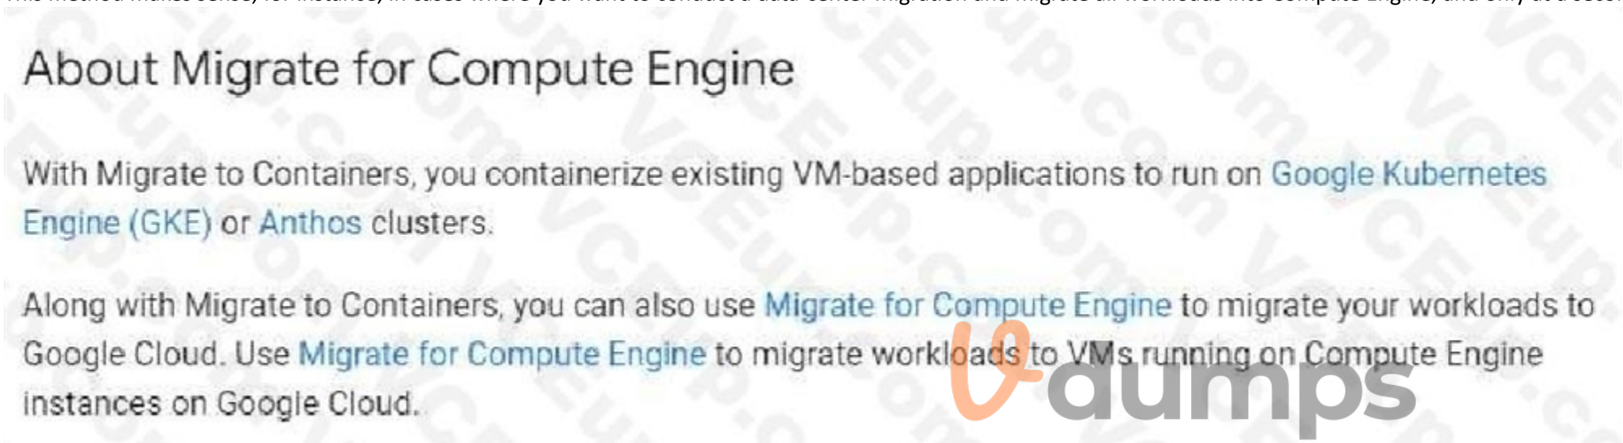


https://cloud.google.com/migrate/containers/docs/architecture

**QUESTION 64**
An IoT platform is providing services to home security systems. They have more than a million customers, each with many home devices. Burglaries or child safety issues are concerns that the clients customers. Therefore, the platform has to respond very quickly in near real time. What could be a typical data pipeline used to support this platform on Google Cloud?

A. Cloud Pub/Sub, Cloud Dataflow, Data Studio
B. Cloud Functions, Cloud Dataproc, Looker
C. Cloud Pub/Sub, Cloud Dataflow, BigQuery
D. Cloud Functions, Cloud Dataproc, BigQuery

**Correct Answer: A**
**Section:**
**Explanation:**
=> Cloud Pub/Sub- Cloud Pub/Sub is the best to be the end-point for ingesting large amounts of data.
It will grow as required, can stream data to downstream systems, and can also work with intermittently available backends.
=> Cloud Dataflow- supports streaming data and therefore is an appropriate option for processing the data that is ingested.
=> BigQuery- BigQuery also supports streaming data and its possible to do real time ana-lytics on it.
=> DataStudio- DataStudio and Looker are for visualization. They don't have any in-built analysis.
=> Cloud Functions- Cloud Functions is a useful serverless endpoint. However, Pub/Sub is better in this case because it can also retain messages for a set period if it was not possi-ble to deliver it first time.
=>Cloud Dataproc- Cloud Dataproc is used for Hadoop/Spark workloads and won't be a good fit here.

**QUESTION 65**
The CFO is attending one of the preliminary meetings in the migration strategy meeting. She brings up the concern about costs. They have contracts with their vendors and the payments they will need to make when purchasing any kind of infrastructure. This gives them a clear view of numbers for resource budgeting and planning. Can she get the same kind of clarity on Google Cloud?

A. Yes. Do a trial run of typical workloads. See the billing amount and that becomes the base reference.
B. Yes, the Cloud Native Computing Foundation publishes yearly numbers on the cost of running the cloud. Use that as a reference.
C. Yes, the Pricing Calculator can be used to estimate the cost of resources.
D. Yes, Google provides a typical cost of application workloads by region and indus-try. Use that as a reference.

**Correct Answer: C**
**Section:**
**Explanation:**
The pricing calculator can be used to give clear estimates of resource usage.
-> Running test loads is as closely indicative as using the pricing calculator.
-> There are no cloud cost references published, either by Google or CNCF. Even if some companies have published such info. It might not apply to you.
Reference link:- https://cloud.google.com/products/calculator

**QUESTION 66**
The government has ordered an audit of your company's dat a. You have hired an external company to conduct the audit. They need to be able to review the data stored in your Cloud Storage buckets across eight projects. How would you grant them access?

A. Give the auditors an Owner role on the eight buckets so that they have proper access.
B. Give them Storage Object Viewer access to the buckets in those eight projects.
C. They might need access to all projects as the audit progresses; so give them access to all Storage buckets so that you don't have to do it repeatedly later on.
D. They might need access to all projects as the audit progresses; so give them the Editor role on all Storage buckets so that you don't have to do it repeatedly later on.

**Correct Answer: B**
**Section:**
**Explanation:**
Apply the Principle of Least Privilege and only provide read permissions on only the required buckets. No more, no less
https://cloud.google.com/storage/docs/access-control/iam-roles

**QUESTION 67**
A prospect wants to be able to store and analyze dat a. Their analysts already know SQL, but are not familiar with other technologies. Which of these databases can the analysts use without addi-tional training?

A. Cloud SQL, BigQuery, Datastore
B. Spanner, Cloud SQL, BigQuery
C. Cloud SQL, Firestore, Datastore
D. Cloud SQL, Bigtable, BigQuery

**Correct Answer: B**
**Section:**
**Explanation:**
Spanner, Cloud SQL, BigQuery
Spanner- Cloud Spanner is a fully managed, mission-critical, relational database service that offers transactional consistency at global scale, automatic, synchronous replication for high availability, and support for two SQL Google Standard SQL and PostgreSQL.
Cloud SQL- Cloud SQL is a fully-managed database service that helps you set up, maintain, manage, and administer your relational databases on Google Cloud Platform.
BigQuery- Google BigQuery is a cloud-based Architecture and provides exceptional performance as it can auto-scale up and down based on the data load and performs data analysis efficiently. On the other hand, SQL Server

is based on client-server architecture and has fixed performance throughout unless the user scales it manually.

**QUESTION 68**
Which of the following is/are true about Bare Metal Solutions?

A. Enterprise-grade deployment platform
B. All your existing investment in tooling and best practices will work as is
C. Continue to run any version, and feature set, any database option, and any cus-tomizations (patchsets)
D. All of the Above.

**Correct Answer: D**
**Section:**
**Explanation:**
Bare Metal Solution for Oracle
Bring your Oracle workloads to Google Cloud with Bare Metal Solution and jumpstart your cloud journey with minimal risk.
- Continue to run any version, any feature set, any database option, and any customizations (patchsets)
- Enterprise-grade deployment platform
- High availability with Oracle RAC
- Works with any application, any Oracle versions
- All your existing investment in tooling and best practices will work as is

**QUESTION 69**
A customer has new applications to build that has to handle both batch data and streaming dat a. Which product should they choose?

A. Dataprep
B. Dataflow
C. Dataproc
D. Data Fusion

**Correct Answer: B**
**Section:**
**Explanation:**
Dataflow is the managed version of Apache Beam. Beam = Batch + Stream. Unified stream and batch data processing that's serverless, fast, and cost-effective.


Reference link- https://cloud.google.com/dataflow

**QUESTION 70**

Your application has repeated data requests of the exact same nature. At the same time, the number of user requests is increasing. Monitoring indicates that the load on the existing database is increasing, and there seems to be a bottleneck. An analysis of the data requested shows us that it is application-managed data and that it changes, but not often. How can you improve the efficiency of the application?

A. Use Cloud Memorystore to improve speed via caching
B. Increase the amount of RAM on the machine hosting the database so that it has higher data throughput.
C. Use Cloud Storage with multi-regional storage so that all users accessing the data will have lower latency
D. Increase the number of CPUs on the machine hosting the database so that it has higher data throughput.

**Correct Answer: A**
**Section:**
**Explanation:**
Cloud Memorystore is an in-memory database that has sub-millisecond latency. This is ideal for caching application data that also changes once in a while.
https://cloud.google.com/memorystore

**QUESTION 71**
What conditions be true if a VM interface wants to send packets to the external IP addresses of Google APIs and services using Private Google Access?

A. VM interface does not have an external IP address assigned.
B. VM interface is connected to a subnet where Private Google Access is disabled
C. Both A and B
D. None of the Above.

**Correct Answer: A**
**Section:**
**Explanation:**
A VM interface can send packets to the external IP addresses of Google APIs and services using Private Google Access if all these conditions are met:
- The VM interface is connected to a subnet where Private Google Access is enabled.
- The VPC network that contains the subnet meets the network requirements for Google APIs and services.
- The VM interface does not have an external IP address assigned.
- The source IP address of packets sent from the VM matches the VM interface's primary internal IP address or an internal IP address from an alias IP range.
A VM with an external IP address assigned to its network interface doesn't need Private Google Access to connect to Google APIs and services. However, the VPC network must meet the requirements for accessing Google APIs and services.

**QUESTION 72**
Your customer has reliable information to indicate that they will use a certain amount of computing and analytics. The workloads are critical and they don't want to take a chance with VMs or BigQuery slots being unavailable during a peak period. How can they ensure that they allocate the capacity?

A. Send in the filled form to Google Cloud support to reserve the Compute Engine and BigQuery resources required.
B. Create reservations on Compute Engine and BigQuery.
C. On the day the capacity is required, set a scheduled job that will provision as many resources as required and lock it in.
D. Google Cloud is elastic for resources. You cannot reserve resources in advance; it is pay per use.

**Correct Answer: B**
**Section:**
**Explanation:**
Create reservations on Compute Engine and BigQuery. You can reserve capacity in advance and use it over a period of time. You could also get a cost advantage.
=> There is no need for involved support. It is self-serve via the console.
=> You can reserve resources in advance when you have the need for it. And when you want to take a pay-per-use approach, that is also possible.
=> It is not a good idea to be lock in/hoard resources; you'll pay unnecessarily for resources. Also, it is difficult to time exactly when the demand will be.
References:
https://cloud.google.com/compute/docs/instances/reserving-zonal-resources

https://cloud.google.com/bigquery/docs/reservations-intro

**QUESTION 73**
An organization's applications run on an inflexible, on-premises architecture. The organization has decided to modernize their existing applications with the cloud. What may have prompted this business decision?

A. Developers want cloud providers to take full control of their application performance.
B. IT managers want cloud providers to automatically deploy their infrastructure.
C. IT managers want to stop making gradual changes.
D. Developers want to test ideas and experiment with more ease.

**Correct Answer: D**
**Section:**
**Explanation:**
Modernizing applications means they can make alterations and innovate more easily.

**QUESTION 74**
An organization wants to scale their existing virtual machine architecture as quickly as possible. Why should the organization use VMware Engine?

A. To archive virtual machine instances.
B. To deploy custom APIs seamlessly.
C. To migrate virtual machines to containers.
D. To replatform virtual machines as they are.

**Correct Answer: D**
**Section:**
**Explanation:**
VMware Engine helps migrate and run virtual machines in Google Cloud with minimal changes to the VM architecture.

A virtual machine (VM) is a digital version of a physical computer. Virtual machine software can run programs and operating systems, store data, connect to networks, and do other computing functions, and requires maintenance such as updates and system monitoring. Multiple VMs can be hosted on a single physical machine, often a server, and then managed using virtual machine software. This provides flexibility for compute resources (compute, storage, network) to be distributed among VMs as needed, increasing overall efficiency. This architecture provides the basic building blocks for the advanced virtualized resources we use today, including cloud computing.

Learn about virtual machines and VM family types that are available with Compute Engine, the cloud-based computing infrastructure from Google Cloud.

https://cloud.google.com/learn/what-is-a-virtual-machine

**QUESTION 75**
Your Google Cloud Platform [GCP] admin has to manage a bunch of API keys for external services that are accessed by different applications, which are used by a few teams. What is the best way to manage them?

A. Share the information in a Github repository and grant access to the repo in IAM as required.
B. Store the information in Secret Manager and give IAM read permissions as re-quired.
C. Store the information in Kubernetes Secrets and only grant read permissions to users as required.
D. Encrypt the information and store it in Cloud Storage for centralized access. Give the decrypt key only to the users who need to access it.

**Correct Answer: B**
**Section:**
**Explanation:**
Store the information in Secret Manager is a secure and convenient storage system for API keys, passwords, certificates, and other sensitive data. Secret Manager provides a central place and single source of truth to manage access, and audit secrets across Google Cloud.
https://cloud.google.com/secret-manager

**QUESTION 76**
What are the key features of Google Cloud Identity.

A. Multi-factor authentication (MFA)
B. Single sign-on (SSO)
C. Works with your favorite apps and Endpoint management
D. All of the Above

**Correct Answer: D**
**Section:**
**Explanation:**
Cloud Identity:
A unified identity, access, app, and endpoint management (IAM/EMM) platform.
- Give users easy access to apps with single sign-on.
- Multi-factor authentication protects user and company data.
- Endpoint management enforces policies for personal and corporate devices KEY FEATURES :
Modernize IT and strengthen security
Multi-factor authentication (MFA)
Help protect your user accounts and company data with a wide variety of MFA verification methods such as push notifications, Google Authenticator, phishing-resistant Titan Security Keys, and using your Android or iOS device as a security key.
Endpoint management
Improve your company's device security posture on Android, iOS, and Windows devices using a unified console. Set up devices in minutes and keep your company data more secure with endpoint management. Enforce security policies, wipe company data, deploy apps, view reports, and export details.
Single sign-on (SSO)
Enable employees to work from virtually anywhere, on any device, with single sign-on to thousands of pre-integrated apps, both in the cloud and on-premises.
Works with your favorite apps
Cloud Identity integrates with hundreds of cloud applications out of the boxóand we're constantly adding more to the list so you can count on us to be your single identity platform today and in the future.

**QUESTION 77**
A partner of yours used to have their own private data center. Your company was already on Google Cloud and now they have also moved to Google Cloud. You are investigating whether there are ways to collaborate better or shared services. What would be one good option to consider?

A. Use Private Service Access within Google Cloud.
B. Use VPC Peering to share resources privately between your two organizations.
C. Use public IP addresses as before. It will automatically be routed internally only.
D. Use VPC Shared Networks to share common resources.

**Correct Answer: B**
**Section:**
**Explanation:**
VPC Network Peering allows internal IP address connectivity across two Virtual Private Cloud (VPC) networks regardless of whether they belong to the same project or the same organization.
-> Shared VPC is only within an organization - it allows an organization to connect resources from multiple projects to a common Virtual Private Cloud (VPC) network, so that they can communicate with each other securely and efficiently using internal IPs from that network.
-> Private Google Access is only to access Google APIs and services
References:
-> https://cloud.google.com/vpc/docs/vpc-peering
-> https://cloud.google.com/vpc/docs/private-google-access
-> https://cloud.google.com/vpc/docs/shared-vpc

**QUESTION 78**
What are the network requirements for Private Google Access?

A. Private Google Access automatically enables any API.
B. Your network must have appropriate routes for the destination IP ranges used by Google APIs and services.
C. Both A and B
D. None of the Above

**Correct Answer: B**
**Section:**
**Explanation:**
Network requirements for Private Google Access:
- Because Private Google Access is enabled on a per-subnet basis, you must use a VPC network. Legacy networks are not supported because they don't support subnets.
- Private Google Access does not automatically enable any API. You must separately enable the Google APIs you need to use via the APIs & services page in the Google Cloud Console.
- If you use the private.googleapis.com or the restricted.googleapis.com domain names, you'll need to create DNS records to direct traffic to the IP addresses associated with those domains.
- Your network must have appropriate routes for the destination IP ranges used by Google APIs and services. These routes must use the default internet gateway next hop. If you use the private.googleapis.com or the restricted.googleapis.com domain names, you only need one route (per domain). Otherwise, you'll need to create multiple routes.
- Egress firewalls must permit traffic to the IP address ranges used by Google APIs and services. The implied allow egress firewall rule satisfies this requirement. For other ways to meet the firewall requirement.

**QUESTION 79**
A fitness band company is continuously ingesting data from millions of its consumers. Different kinds of data based on time, like location, heartbeat rate, temperature, movement, etc. are connect-ed.
They need a high throughput database that can write data very fast. Since their users are spread across the world, they need the database to be geographically scalable. Consumers also want to see near-real-time visualizations of their activities. Which of these databases would be a good fit?

A. Cloud SQL
B. Bigtable
C. Spanner
D. Firestore

**Correct Answer: B**
**Section:**
**Explanation:**
Bigtable is the best suited for time series data. It also has high read-write throughput and ability to scale globally.

**QUESTION 80**
Your team is working on building a machine learning model. There are a bunch of terminologies that are being used. What is an "instance" or an "example"?

A. An input variable is used in making predictions. E.g. number of rooms in a house price prediction model.
B. One row of a dataset containing one or more input columns and possibly a prediction result.
C. An answer for a prediction task,¨ either the answer produced by a machine learning system or the right answer supplied in training data. E.g. image contains a "cat".
D. The "knobs" that you tweak during successive runs of training a model. E.g. learning rate

**Correct Answer: B**
**Section:**
**Explanation:**
One row of a dataset containing one or more input columns and possibly a prediction result.

- **Instance**: The thing about which you want to make a prediction. For example, the instance might be a web page that you want to classify as either "about cats" or "not about cats".
- **Label**: An answer for a prediction task  either the answer produced by a machine learning system, or the right answer supplied in training data. For example, the label for a web page might be "about cats".
- **Feature**: A property of an instance used in a prediction task. For example, a web page might have a feature "contains the word 'cat'".
- **Feature Column**: A set of related features, such as the set of all possible countries in which users might live. An example may have one or more features present in a feature column. "Feature column" is Google-specific terminology. A feature column is referred to as a "namespace" in the VW system (at Yahoo/Microsoft), or a field.
- **Example**: An instance (with its features) and a label.
- **Model**: A statistical representation of a prediction task. You train a model on examples then use the model to make predictions.

https://developers.google.com/machine-learning/guides/rules-of-ml#terminology

**QUESTION 81**
A retail store has discovered a cost-effective solution for creating self-service kiosks. They can use existing check-out hardware and purchase a virtual customer service application. Why do they also need an API?

A. To connect the check-out hardware to the public cloud.
B. To connect the new application with the legacy system.
C. To migrate all customer data for disaster recovery.
D. To update the check-out hardware remotely.

**Correct Answer: B**
**Section:**
**Explanation:**
APIs can create new business value by connecting legacy systems (the checkout hardware) with new software (the virtual customer service application).

**QUESTION 82**
Your customer is making a decision on whether to move to Google Cloud. Their key concern is about 10,000 VMs that are part of their IT infrastructure used across more than 110 applications. They are apprehensive of too many changes at this stage. They want to get to Google Cloud in the easiest way possible with minimal disruption. What option would you recommend for them?

A. Use Migrate for Anthos
B. Lift and shift the VMs to serverless options like App Engine Flex.
C. Re-architect on-prem to use Kubernetes and then slowly extend and bridge the on-prem data center to the Google Cloud data center.
D. Use Migrate for Compute

**Correct Answer: D**
**Section:**
**Explanation:**
Migrate for Compute Engine's advanced replication migration technology copies instance data to Google Cloud in the background with no interruptions to the source workload that's running.

Cloud migration creates a lot of questions. Migrate for Compute Engine by Google Cloud has the answers. Whether you're looking to migrate one application from on-premises or one thousand enterprise-grade applications across multiple data centers, Migrate for Compute Engine gives any IT team, large or small, the power to migrate their workloads to Google Cloud.

Watch the video to your right to hear what one of our customers, Rackspace Technology, thinks about Migrate for Compute Engine's speed and ease of use.

https://cloud.google.com/migrate/compute-engine

**QUESTION 83**
You are leading projects in an IT services company. Your customer's project requires analyzing images.
They have many 10s of 1000s of raw images that they have made available to you. Your small technology team needs to build a machine learning model. The images are unlabeled. You don't have the people or the capacity to label the images. What is your approach?

A. Look for open-source labeled images that closely resemble the given images.
B. Request data labeling service from Google.
C. Tell the customer it is their duty to label the images.
D. Hire temporary workers who can quickly label the images.

**Correct Answer: C**
**Section:**
**Explanation:**
Google's Data Labeling Service lets you work with human labelers to generate highly accurate labels for a collection of data that you can use in machine learning models.
References:
-> https://cloud.google.com/vertex-ai/docs/datasets/data-labeling-job-> https://cloud.google.com/ai-platform/data-labeling/docs

**QUESTION 84**
You are working with the head of the IT team and planning the move of computing systems. The questionnaire indicates that they have a reporting application that runs almost 24 hours every day of the week. When there is extra load, it queues up the processing and executes tasks when there is less demand. Which of these compute options would you recommend for them?

A. Use a serverless option - App Engine Standard for Flex
B. Use a server-based option - Compute Engine.
C. Use a serverless option - Cloud Functions
D. Serverless option - Cloud Run

**Correct Answer: C**
**Section:**
**Explanation:**
- Because Compute Engine VMs are the preferred compute option as they are long-running.

**QUESTION 85**
With respect to the Core Feature of Standby Instances of Cloud SQL which one of the options is correct.?

A. The standby instance is used in high availability to replace the primary instance when failover occurs. The standby instance appears in the Google Cloud Console but does not get billed. When failover occurs, connections to the primary instance need to be manually transferred to the standby instance.
B. The standby instance is used in high availability to replace the primary instance when failover occurs. The standby instance appears in the Google Cloud Console but does not get billed. When failover occurs, connections to the primary instance are automatically transferred to the standby instance.
C. The standby instance is used in high availability to replace the primary instance when failover occurs. The standby instance doesn't appear in the Google Cloud Console. When failover occurs, connections to the primary instance are automati-cally transferred to the standby instance.
D. None of the Above.

**Correct Answer: C**
**Section:**
**Explanation:**
The standby instance is used in high availability to replace the primary instance when failover occurs.
The standby instance doesn't appear in the Google Cloud Console. When failover occurs, connections to the primary instance are automatically transferred to the standby instance.
Cloud SQL Key Terms:
Cloud SQL instance
A Cloud SQL instance corresponds to one virtual machine (VM). The VM includes the database instance and accompanying software containers to keep the database instance up and running.
Database instance
A database instance is the set of software and files that operate the databases: MySQL, PostgreSQL or SQL Server.
High availability
Cloud SQL instances using high availability (HA) provide greater reliability than non-HA instances.
HA in Cloud SQL works by having two synchronized instances: a primary instance and a standby instance. Each instance has exactly one VM. Each instance is in a different zone in the same region.
Failover
A failover is when Cloud SQL switches serving from the original primary instance to the standby instance.
Autofailover is a mechanism that automatically triggers failover when a Cloud SQL instance didn't issue a heartbeat in the previous interval.
Standby instances
The standby instance is used in high availability to replace the primary instance when failover occurs.
The standby instance doesn't appear in the Google Cloud Console. When failover occurs, connections to the primary instance are automatically transferred to the standby instance.
Clone
When you clone a Cloud SQL instance, you create a new instance that is a copy of the source instance, but is completely independent. After cloning is complete, changes to the source instance are not reflected in the clone, and changes in the clone are not reflected in the source instance.
Replication
Replication is the ability to create copies of a Cloud SQL instance or an on-premises database, and offload work to the copies. The main reason for using replication is to scale the use of data in a database without degrading performance on the primary instance.
Read replica
The read replica is an exact copy of the primary instance. Data and other changes on the primary instance are updated in almost real time on the read replica. Send your write transactions to the primary instance, and your read requests to the read replica. The read replica processes queries, read requests, and analytics traffic, thus reducing the load on the primary instance.

Source server
Replication copies transactions from a primary instance to one or more read replicas. The primary instance is also called the source server. The source server can be a Cloud SQL primary instance, or a server outside of Google Cloud, such as an on-premises server or a server running in a different cloud. If the source server is outside of Google Cloud, we call it Replication from an external server.
Cloud SQL Auth proxy client
The Cloud SQL Auth proxy client is open source software maintained by Cloud SQL. It connects to a companion process, the Cloud SQL Auth proxy server, running on your Cloud SQL instance. You run the Cloud SQL Auth proxy client on your own servers. The Cloud SQL Auth proxy client can be used to establish a secure SSL/TLS connection to the database instance, and/or to avoid having to open the firewall. Authentication is done through Identity and Access
Management (IAM).

**QUESTION 86**
You are a database manager working for a new product that will need millions of reading and writing from the database, with zero downtime, key-value i.e. NoSQL features, no manual steps should be required to ensure consistency, repair data, synchronize writes and deletes, Which of the following database you choose?

A. Cloud SQL
B. Cloud BigTable
C. Cloud Spanner
D. Cloud Firestore

**Correct Answer: B**
**Section:**
**Explanation:**
Cloud BigTable
Key features
High throughput at low latency
Bigtable is ideal for storing very large amounts of data in a key-value store and supports high read and write throughput at low latency for fast access to large amounts of data. Throughput scales linearlyóyou can increase QPS (queries per second) by adding Bigtable nodes. Bigtable is built with proven infrastructure that powers Google products used by billions such as Search and Maps.
Cluster resizing without downtime
Scale seamlessly from thousands to millions of reads/writes per second. Bigtable throughput can be dynamically adjusted by adding or removing cluster nodes without restarting, meaning you can increase the size of a Bigtable cluster for a few hours to handle a large load, then reduce the cluster's size againóall without any downtime.
Flexible, automated replication to optimize any workload
Write data once and automatically replicate where needed with eventual consistencyógiving you control for high availability and isolation of reading and write workloads. No manual steps are needed to ensure consistency, repair data, or synchronize writes and deletes. Benefit from a high availability SLA of 99.999% for instances with multi-cluster routing across 3 or more regions (99.9% for single-cluster instances).

**QUESTION 87**
In terms of Cloud SQL for MySQL Features offered by Google Cloud Platform which of the statements is/are correct?

A. Do not support Private IP (private service access).
B. Customer data is encrypted on Google's internal networks and in database tables, temporary files, and backups.
C. Do not Provide automated and on-demand backups and point-in-time recovery.
D. None of the above

**Correct Answer: B**
**Section:**
**Explanation:**
Cloud SQL for MySQL:
Features
- Fully managed MySQL Community Edition databases in the cloud.
- Cloud SQL instances support MySQL 8.0, 5.7 (default), and 5.6, and provide up to 624 GB of RAM and 64 TB of data storage, with the option to automatically increase the storage size, as needed.
- Create and manage instances in the Google Cloud Console.

- Instances are available in the Americas, EU, Asia, and Australia.
- Customer data is encrypted on Google's internal networks and in database tables, temporary files, and backups.
- Support for secure external connections with the Cloud SQL Auth proxy or with the SSL/TLS protocol.
- Support for private IP (private services access).
- Data replication between multiple zones with automatic failover.
- Import and export databases using mysqldump, or import and export CSV files.
- Support for MySQL wire protocol and standard MySQL connectors.
- Automated and on-demand backups and point-in-time recovery.
- Instance cloning.
- Integration with Google Cloud's operations suite logging and monitoring.

**QUESTION 88**
What issues can arise when organizations integrate third-party systems into their cloud infrastructure?

A. Third-party systems may not be powerful enough to run many critical business applications.
B. Without sufficient security measures and regular checks, unsecured third-party systems can pose a threat to data security.
C. Over-reliance on third-party systems limits an organization's potential for innova-tion.
D. Third-party systems are less capable of addressing an organization's security re-quirements.

**Correct Answer: B**
**Section:**
**Explanation:**
Because unsecured third-party systems are a cybersecurity threat.

**QUESTION 89**
A customer of yours has an SLA with their client that a particular service will respond within 4 seconds.
The end client has reported that it feels slower. Your engineers do a trial at the client site and notice that there seems to be a delay for many of the requests. It's your team's responsibility to identify the issue quickly within the strict timeline for fixes according to the contract, and then fix it. What should you do?

A. Recommend a move to serverless technologies which will scale automatically on demand.
B. Add logging statements at multiple points in the application, build it, and deploy it. Now new requests will give us information on latency in the logs.
C. Check if the browsers used by the client are different from yours. If they are, that's most likely the issue. Ensure that everybody uses the latest version of the browser that you are also using.
D. Use Cloud Trace to collect latency data and track how requests propagate and why there is a delay.

**Correct Answer: D**
**Section:**
**Explanation:**
Cloud Trace is a built-in tool in the Operations suite to identify issues like latency.
-> Such fixes are unlikely to change core issues like the service itself being architected or written suboptimally.
Though changes like browser, networking, etc. are helpful, it would be the wrong approach to first recommend that the customer upgrade all their hardware and software.
-> Rewriting code and logging information is going to be time consuming. In general though, logging should always be included in code and it can give good insights. But tracing is way more specific and comprehensive for this requirement.
-> In certain cases, we might identify scaling as the issue. But we should first identify the core problem. So, start with tracing. We can also achieve scale in server-ful technologies.
Reference link- https://cloud.google.com/trace

**QUESTION 90**
A large organization is struggling to manage their cloud costs effectively. They want to increase visibility into cloud costs. Which cost management approach should the organization use?

A. Establish a partnership between finance, technology, and business teams.
B. Appoint a single person to monitor cloud spending across the organization.
C. Review any cloud spending that exceeds the organization's error budget.
D. Increase monitoring of on-premises infrastructure and services.

**Correct Answer: A**
**Section:**
**Explanation:**
Because cross-team partnerships are part of the visibility cost management strategy.
https://wa.aws.amazon.com/wat.question.COST_1.en.html

**QUESTION 91**
A financial services company is running an experimental application workload that has a very large number of mathematical calculations involving floating-point numbers. The current application that is running on compute engine is not providing enough speed and throughput. What are the options to increase the processing performance?

A. Use a serverless option like Cloud Functions that will automatically scale as much as required.
B. Instead of using a "general purpose" machine family, use "compute-optimized" machine family.
C. Since processing could also be dependent on reading and writing data to the disk, use a fast Local SSD.
D. Attach GPUs to the virtual machine for number crunching.

**Correct Answer: D**
**Section:**
**Explanation:**
Compute Engine provides graphics processing units (GPUs) that you can add to your virtual machines (VMs). You can use these GPUs to accelerate specific workloads on your VMs such as machine learning and data processing.
https://cloud.google.com/compute/docs/gpus

**QUESTION 92**
Your client is a financial services company giving loans based on customer profiles. As part of the regulatory compliance, they have to collect a bunch of different documents with know your customer (KYC) information. They want to be able to process the information in these documents quickly and at scale. They want to integrate the chosen solution as quickly as possible. What are your options on Google Cloud?

A. Integrate the Cloud Vision API to create a custom model to handle the documents.
B. Create a model using TensorFlow and integrated it into the process workflow.
C. Integrate the Lending DocAI and Document AI in two there processes workflow of the processing loan requests.
D. Integrate the Natural Language API to read the request sent in by clients and to process the forms.

**Correct Answer: C**
**Section:**
**Explanation:**
Lending DocAI is a pre-packaged AI solution that speeds "up the mortgage workflow processes to easily process loans and automate document data capture, while ensuring the accuracy and breadth of different documents (e.g., tax statements and asset documents)."
https://cloud.google.com/solutions/lending-doc-ai

**QUESTION 93**
A large travel company has thus far invested heavily in their technology team. There is strategic pressure on the company to focus on their core business and innovate to survive in certain geographies and thrive in others. They are evaluating whether a move to Google Cloud will be good for them. Which of these reasons would be relevant for them? (choose two answer)

A. Application architecture won't be too involved because of serverless options.

B. The IT team won't have to manage software upgrades, security patches, etc. for the VMs.
C. The IT team won't have to work on procuring and provisioning new hardware and refreshes to existing hardware.
D. Budgeting won't be an issue since the cloud takes care of billing.

**Correct Answer: B, C**
**Section:**

**QUESTION 94**
While on-premise, an enterprise had multiple teams, each with its own analytics data store.
Attempts to converge the storage for centralized, company-wide analysis failed because of speed and scaling issues. What would be the preferred destination architecture on Google Cloud?

A. Migrate to Bigtable which provides high throughput reads and writes.
B. Migrate to Cloud Spanner as a globally scalable SQL database.
C. Migrate to BigQuery as a central data warehouse.
D. Migrate to Cloud SQL which supports multiple databases like MySQL, PostgreSQL, and SQL Server - all of the customer's SQL databases can be accommodated here.

**Correct Answer: C**
**Section:**
**Explanation:**
BigQuery is the data warehousing option on Google Cloud. Since the source data has already been used for analysis, it should easily fit the BigQuery structure too.

**QUESTION 95**
Keeping Flavours of Apigee in mind, which of the following statements is/are correct?

A. A hybrid version consisting of a runtime plane installed on-premises or in a cloud provider of your choice, and a management plane running in Apigee's cloud. In this model, API traffic and data are confined within your own enterprise-approved boundaries.
B. A hosted SaaS version in which Apigee maintains the environment, allowing you to concentrate on building your services and defining the APIs to those services.
C. There are two types of Flavours in Apigee i.e. Apigee & Apigee Hybrid.
D. All of the above are correct.

**Correct Answer: D**
**Section:**
**Explanation:**
Flavors of Apigee
Apigee comes in the following flavors:
Apigee: A hosted SaaS version in which Apigee maintains the environment, allowing you to concentrate on building your services and defining the APIs to those services.
Apigee hybrid: A hybrid version consisting of a runtime plane installed on-premises or in a cloud provider of your choice, and a management plane running in Apigee's cloud. In this model, API traffic and data are confined within your own enterprise-approved boundaries.

**QUESTION 96**
Cloud SQL is a fully-managed relational database service for MySQL, PostgreSQL and SQL serv-ers, keeping Cloud SQL Google Cloud Service in mind, which of the following statements is/are correct?

A. Data inside cloud SQL is automatically Encrypted.
B. Cloud SQL automatically ensures your databases are reliable, secure, and scala-ble so that your business continues to run without disruption.
C. With DMS (Database Migration Service) it becomes very easy to Migration of Pro-duction Database.
D. All of the above

**Correct Answer: D**
**Section:**
**Explanation:**
Cloud SQL
Fully managed relational database service for MySQL, PostgreSQL, and SQL Server. Run the exact same relational databases you know with their rich extension collections, configuration flags and developer ecosystem, but without the hassle of self management.
- Reduce maintenance cost with fully managed MySQL, PostgreSQL and SQL Server databases.
- Ensure business continuity with reliable and secure services backed by 24/7 SRE team.
- Automate database provisioning, storage capacity management, and other time-consuming tasks.
- Database observability made easy for developers with Cloud SQL Insights.
- Easy integration with existing apps and Google Cloud services like GKE and BigQuery.
Key features:
Fully managed
Cloud SQL automatically ensures your databases are reliable, secure, and scalable so that your business continues to run without disruption. Cloud SQL automates all your backups, replication, encryption patches, and capacity increasesó while ensuring greater than 99.95% availability, anywhere in the world.
Integrated
Access Cloud SQL instances from just about any application. Easily connect from App Engine, Compute Engine, Google Kubernetes Engine, and your workstation. Open up analytics possibilities by using BigQuery to directly query your Cloud SQL databases.
Reliable
Easily configure replication and backups to protect your data. Go further by enabling automatic failover to make your database highly available. Your data is automatically encrypted, and Cloud SQL is SSAE 16, ISO 27001, and PCI DSS compliant and supports HIPAA compliance.
Easy migrations to Cloud SQL
Database Migration Service (DMS) makes it easy to migrate your production databases to Cloud SQL with minimal downtime. This serverless offering eliminates the manual hassle of provisioning, managing, and monitoring migration-specific resources. DMS leverages the native replication capabilities of MySQL and PostgreSQL to maximize the fidelity and reliability of your migration. And it's available at no additional charge for native like-to-like migrations to Cloud SQL.

**QUESTION 97**
Customer Managed Encryption Keys (CMEK) can be used for encrypting data inside Cloud BigTable, which of the following statements is/are correct. (Select two answer)

A. Administrators can not rotate
B. Not supported for instances that have clustered in more than one region.
C. CMEK can only be configured at the cluster level.
D. You can not use the same CMEK key in multiple projects

**Correct Answer: B, C**
**Section:**
**Explanation:**
Customer-managed encryption keys for Cloud BigTable.
By default, all the data at rest in Cloud Bigtable is encrypted using Google's default encryption.
Bigtable handles and manages this encryption for you without any additional action on your part.
If you have specific compliance or regulatory requirements related to the keys that protect your data, you can use customer-managed encryption keys (CMEK) for BigTable. Instead of Google managing the encryption keys that protect your data, your BigTable instance is protected using a key that you control and manage in Cloud Key Management Service (Cloud KMS).
Features
Security: CMEK provides the same level of security as Google's default encryption but provides more administrative control.
Data access control: Administrators can rotate, manage access to, and disable or destroy the key used to protect data at rest in BigTable .
Auditability: All actions on your CMEK keys are logged and viewable in Cloud Logging.
Comparable performance: BigTable CMEK-protected instances offer comparable performance to BigTable instances that use Google default encryption.
Flexibility: You can use the same CMEK key in multiple projects or instances or you can use separate keys, depending on your business needs.

**QUESTION 98**
You have contracted a partner to conduct some medical trials. This is a limited, 2-month contract. At the end of each day, you are expecting about 10 Gbs of dat a. The data is highly sensitive. What networking option would you employ?

A. As the name indicates, set up Partner Interconnect with your partner company.
B. Setup Dedicated Interconnect with your partner.
C. Setup Cloud VPN and create an IPsec VPN tunnel with your partner.
D. Create a public IP for a VM and share that with your partners so that they can access it over the internet and share the data.

**Correct Answer: C**
**Section:**
**Explanation:**
"Cloud VPN securely extends your peer network to Google's network through an IPsec VPN tunnel.
Traffic is encrypted and travels between the two networks over the public internet. Cloud VPN is useful for low-volume data connections. For additional connection options, see the Hybrid Connectivity product page."

**QUESTION 99**
You are working with a government agency. A web application serves users of the country. It al-lows citizens to receive certain services in providing their national identity. Citizens have com-plained that they are seeing delays in web page loading compared to before. On investigating, they are seeing a lot of spurious traffic coming in from a few IPs which they have identified as for-eign. What should they do?

A. Setup Firewall rules to deny access to the malicious IPs.
B. Setup Cloud Armor and add the malicious IPs to the deny list.
C. Setup Firewall rules to allow access only to the IPs from within the country.
D. Setup Cloud NAT and remove all the internal IPs and replace it with a single public IP.

**Correct Answer: B**
**Section:**
**Explanation:**
Cloud Armor provides DDoS protection for applications. It can also "Filter your incoming traffic based on IPv4 and IPv6 addresses or CIDRs. Enforce geography-based access controls to allow or deny traffic based on source geo using Google's geoIP mapping."

**QUESTION 100**
You are consulting for a client who is migrating to Google Cloud. They presently have a matrix organization.
Their IT environments were managed around projects. Each team had multiple projects.
All the projects had a flat structure under the company. What would you advise them when plan-ning for the move?

A. On Google Cloud, create a folder corresponding to each team. Under that, there could be projects or further sub folders as the team decides.
B. In terms of not disturbing the project developers and testers, advise them that the strategic decision is to retain the structure on Google Cloud also.
C. Since a Project could spawn other sub-Projects, on Google Cloud it is better to as-sign a folder for each Project.
D. The flat structure is what is currently used in IT organizations, and this can be used as-is which will provide the best results.

**Correct Answer: A**
**Section:**
**Explanation:**
Folders for a related group of projects are the recommended approach.
-> A flat structure under the organization node is possible on Google Cloud, but it is not recommended. It becomes tougher to manage.
-> Projects cannot have sub-projects; there can only be resources within Projects.
Reference link- https://cloud.google.com/resource-manager/docs/cloud-platform-resourcehierarchy

**QUESTION 101**
Your company has a requirement to run manual tests on their web products for UX research before it is released to end customers. The people who will do the tests are external to the company. They will either use their own Gmail id or be given temporary email ids using the applications and record-ing their inputs in another app. The UX testing is done in the last week of the month. Each month the UX testers could be different. How should the IT team manage the users?

A. Since the app is anyways going to be public, create permanent credentials for the UX testers that they can conveniently use each time.
B. It would be a security issue to have users come and go. Recommend that the test-ers be permanently hired to plug the vulnerability issue.
C. It would be a security issue to have users come and go. Recommend that the test-ers be permanently hired to plug the vulnerability issue.
D. Create a Group with the permissions required to do the test and record their in-puts. When users arrive each week, add them to the group and after the testing period, remove them from the group.

**Correct Answer: D**
**Section:**
**Explanation:**
Groups are convenient to use for this requirement. Permissions to the group are automatically inherited by the members of the group. Adding and removing UX testers from the group will grant and remove permissions.

**QUESTION 102**
What is a key difference between VMs and containers?

A. Virtual Machines take less time to launch; containers take longer to launch.
B. Virtual Machines can only run Linux; containers can run any operating system.
C. Virtual Machines use a shared operating system and are therefore lighter; containers are heavier on resources.
D. Each Virtual Machine in a machine has its own operating system; containers will share the same operating system.

**Correct Answer: D**
**Section:**
**Explanation:**
VMs have their individual OSs. All containers on a node use the host operating system.

**QUESTION 103**
In discussions with a prospective customer who wants to move to Google Cloud to make use of the latest, scalable technologies available therein, you learn that there are very strict regulations concern-ing the storage of dat a. They only have the approval to store it in their current private data cen-ter. What would you advise them?

A. Retain on-premise itself those portions of data and compute which are under regulation. Take advantage of all the other cloud capabilities for remaining work-loads.
B. It is too risky to touch anything in such a scenario. It is best to remain entirely on-premise.
C. Regulations are guidelines. As long as the data remains encrypted, you can move it anywhere.
D. Petition the government for changes to such regulations as all industries are mov-ing to the public cloud. Then, when the regulations are eased, move to Google Cloud.

**Correct Answer: A**
**Section:**
**Explanation:**
Moving to Google Cloud is not an all-or-nothing option. Certain workloads can continue to remain on-premise while the predominant chunk moves to Google Cloud

**QUESTION 104**
Which of the followings are core components of Anthos?

A. Infrastructure, container, and cluster management
B. Secure software supply chain

C. Multicluster & Configuration management
D. All of the above are correct.

**Correct Answer: D**
**Section:**
**Explanation:**

| Core Anthos components | Google Cloud | On-premises | Multi-cloud | Attached clusters |
|---|---|---|---|---|
| Infrastructure, container, and cluster management | GKE<br>Multi Cluster Ingress | Anthos clusters on VMware | Anthos clusters on AWS, Anthos clusters on Azure | |
| Multicluster management | Fleets, fleet-enabled components, and Connect | Fleets, fleet-enabled components, and Connect | Fleets, fleet-enabled components, and Connect | Fleets, fleet-enabled components, and Connect |
| Configuration management | Anthos Config Management | Anthos Config Management | Anthos Config Management | Anthos Config Management |
| Migration | Migrate for Anthos and GKE | Migrate for Anthos and GKE | Migrate for Anthos and GKE | |
| Service management | Anthos Service Mesh<br>Anthos Service Mesh dashboards<br>MeshCA certificate authority | Anthos Service Mesh<br>Grafana and Kiali dashboards<br>Istiod certificate authority | Anthos Service Mesh (AWS only) | Anthos Service Mesh |
| Serverless | Cloud Run for Anthos | Cloud Run for Anthos | | |
| Secure software supply chain | Binary Authorization | Binary Authorization (preview) | | |
| Logging and monitoring | Cloud Logging and Cloud Monitoring for system components | Cloud Logging and Cloud Monitoring for system components | | |
| Marketplace | Kubernetes Applications in Cloud Marketplace | Kubernetes Applications in Cloud Marketplace | | |

**QUESTION 105**
When creating machine learning models, a key initial step is to identify the type of model required.
One of these is the classification model. Which of these statements define a classification model?

A. A type of machine learning model for distinguishing among two or more discrete values. E.g. "book", "car".
B. A type of machine learning model is a meta-model maker, which classifies algo-rithms based on the quality of their output.
C. A type of machine learning model that outputs continuous (typically, floating-point) values. E.g. the predicted price of the house is $120,000.
D. A type of classic model approach that is less used today and which has been re-placed by the regression model.

**Correct Answer: A**
**Section:**
**Explanation:**
A classification model classifies the incoming data into one or more discrete classes.

**QUESTION 106**
You are a DevOps Engineer in an E-commerce company that sells products globally, across the countries, Customers buy products, add them to carts or check-in stock from different parts of the world with different

timestamps, you need to choose a database that can scale globally without any hassle and lots of developer support, it should be consistent across regions, can scale horizontally to support enormous user, automatically replicates, shards and even auto transaction pro-cessing.
Which of the following database do you choose?

A. Cloud SQL
B. Cloud Spanner
C. Cloud Firestore.
D. Cloud Storage.

**Correct Answer: B**
**Section:**
**Explanation:**
Cloud Spanner:
Fully managed relational database with unlimited scale, strong consistency, and up to 99.999% availability.
- Get all the benefits of relational semantics and SQL with unlimited scale
- Start at any size and scale with no limits as your needs grow
- Enjoy high availability with zero scheduled downtime and online schema changes
- Deliver high-performance transactions with strong consistency across regions and continents
- Focus on innovation, eliminating manual tasks with capabilities like automatic sharding Automatic sharding Cloud Spanner optimizes performance by automatically sharding the data based on request load and size of the data. As a result, you can spend less time worrying about how to scale your database and instead focus on scaling your business.
Strong transactional consistency
Purpose-built for external, strong, global transactional consistency.
Regional and multi-regional configurations
No matter where your users may be, apps backed by Cloud Spanner can read and write up-to-date strongly consistent data globally Additionally, when running a multi-region instance, your database is able to survive a regional failure, and offers industry-leading 99.999% availability.
Online schema changes with no downtime
Cloud Spanner users can make a schema change, whether it's adding a column or adding an index while serving traffic with zero downtime. Hence you now have the flexibility to adapt your database to your business needs without compromising on the availability of your application.

**QUESTION 107**
A bank wants to track the success of their existing ATM network, which has been modernized with APIs to instantly notify customers about their transfers. What is the benefit of using Apigee to achieve this goal?

A. It has dashboards that chart dimensions and metrics to report on APIs.
B. It replicates banking APIs to create new business value.
C. It measures and tracks their total cost of ownership (TCO).
D. It allows developers to connect the banking APIs with the public cloud.

**Correct Answer: A**
**Section:**
**Explanation:**
Apigee includes analytics services which allow enterprises to report on various aspects of an API.

**QUESTION 108**
The Border Security Agency has hired your software services firm to build an application for them that will collect information about visas stamped on passports. You are given stamped images. You have to find out which country issued the visa and the period of validity. Pull out this data and put it into a database. Which of these applications would be suitable for that?

A. Use Cloud Vision API - write code to identify the text blocks, copy the data, and store it
B. Use TensorFlow - write code that will identify the type of visa and the bounding text blocks. Copy the data and then store it.

C. Use AutoML - upload other images of visas and run the model creation process which will automatically identify the visas
D. Use Data Labeling service - outsource the work of marking and extracting the in-formation to others.

**Correct Answer: A**
**Section:**
**Explanation:**
Cloud Vision API allows you to programmatically identify images, text, etc. in the document. This would be the best option.
https://cloud.google.com/vision

**QUESTION 109**
Which of the following statements is/are true about Google Cloud BigTable?

A. It is not compatible with Hadoop.
B. It Scales from Giga Byte to Peta Byte with No Downtime.
C. It can not be used in Real-time Ad analytics and tracking thousands of IoT Devices Data.
D. It is an enterprise-level Database that offers relational and non-relational features

**Correct Answer: B**
**Section:**
**Explanation:**
Cloud Bigtable
A fully managed, scalable NoSQL database service for large analytical and operational workloads with up to 99.999% availability.
- Consistent sub-10ms latencyóhandle millions of requests per second
- Ideal for use cases such as personalization, ad tech, fintech, digital media, and IoT
- Seamlessly scale to match your storage needs; no downtime during reconfiguration
- Designed with a storage engine for machine learning applications leading to better predictions
- Easily connect to Google Cloud services such as BigQuery or the Apache ecosystem

**QUESTION 110**
You have a well established development and operations team. Your teams were managing the entire software delivery/deployment cycle on-premise. When migrating to the cloud, you want to continue having this approach. Which is the ideal option for you?

A. PaaS - Platform as a Service
B. SaaS - Software as a Service
C. IDaaS - Identity as a Service
D. IaaS - Infrastructure as a Service

**Correct Answer: D**
**Section:**
**Explanation:**
IaaS - you're given virtualized resources like VMs, Storage, Network. It is your responsibility to manage everything beyond that. This would be similar to what the organization had on-premise.

**QUESTION 111**
Which of the following statements is/are true about Cloud Spanner offered by Google Cloud Platform.

A. It can scale horizontally to support additional capacity.
B. It comes with Zero Downtime, No Maintenance windows, and is proven for large and small workloads.
C. You don't need to shard or replicate data.

D. All of the above.

**Correct Answer: D**
**Section:**
**Explanation:**
Cloud Spanner:
Fully managed relational database with unlimited scale, strong consistency, and up to 99.999% availability.
- Get all the benefits of relational semantics and SQL with unlimited scale
- Start at any size and scale with no limits as your needs grow
- Enjoy high availability with zero scheduled downtime and online schema changes
- Deliver high-performance transactions with strong consistency across regions and continents
- Focus on innovation, eliminating manual tasks with capabilities like automatic sharding.

**QUESTION 112**
You are working in a company that provides different services to its customer. Now it also wants to offer some paid API services to its B2B customers for e.g. google provides google maps API, cloud vision API, and language translation
API. You need to figure out the best solution for the service.

A. Java Programming Spring Boot Framework for to solve the problem of APIs man-agement.
B. Cloud Functions with Firestore and payment gateways integration development.
C. Apigee API Management
D. Frontend & Backend Development with NodeJs and angular etc.

**Correct Answer: C**
**Section:**
**Explanation:**
A top-level idea about Apigee API Management and its offered features can help you solve all questions related to Apigee in Cloud Digital Leader Practice Exam.
Apigee is a platform for developing and managing APIs. By fronting services with a proxy layer, Apigee provides an abstraction or facade for your backend service APIs and provides security, rate limiting, quotas, analytics, and more.
Apigee services: The APIs that you use to create, manage, and deploy your API proxies.
Apigee runtime: A set of containerized runtime services in a Kubernetes cluster that Google maintains. All API traffic passes through and is processed by these services.

**QUESTION 113**
Which of these are defined by the following statement: a contract you have with your end customers, which, if you don't meet, you might even have to pay fines?

A. SLA - Service Level Agreement
B. SLC - Service Level Contract
C. SLO - Service Level Objective
D. SLI - Service Level Indicator

**Correct Answer: A**
**Section:**
**Explanation:**

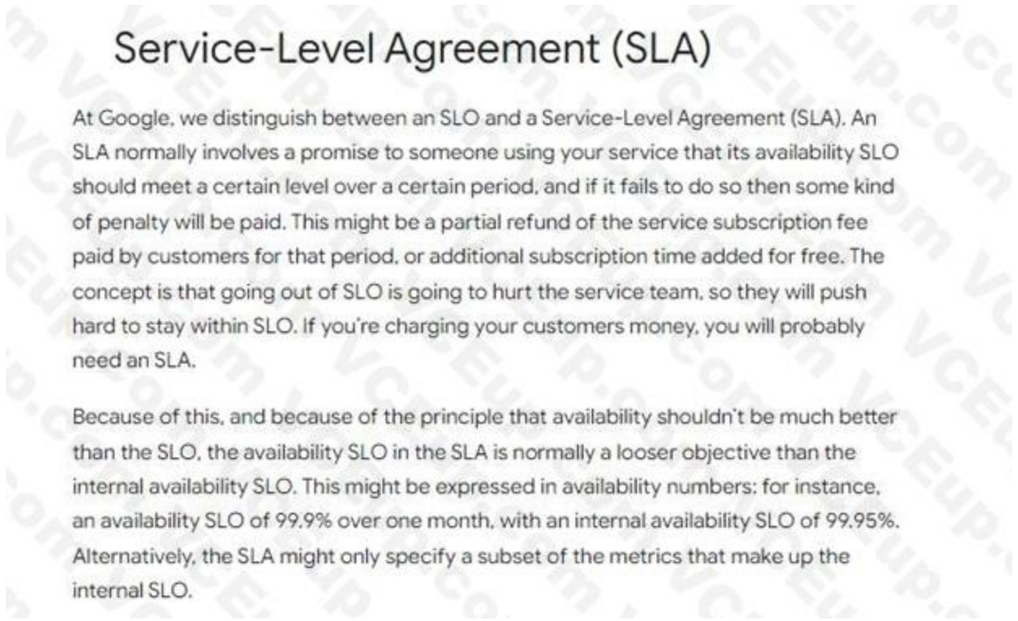

https://cloud.google.com/blog/products/devops-sre/sre-fundamentals-slis-slas-and-slos

**QUESTION 114**
How does a least privilege resource access model contribute to cloud security?

A. Google is responsible for determining access to cloud resources.
B. Employees may only access on-premises software with special permission.
C. Only managers and other senior employees have cloud resource access.
D. Employees only have access to the cloud resources necessary for their job.

**Correct Answer: D**
**Section:**
**Explanation:**
This is the definition of a least privilege model.
A supporting principle that helps organizations achieve these goals is the principle of least privilege.
The principle of least privilege addresses access control and states that an individual should have only the minimum access privileges necessary to perform a specific job or task and nothing more

**QUESTION 115**
You are working for a hospital that stores its medical images in an on-premises data room and it is provided that the hospitals want to use Cloud Storage for archival storage of these images. You are required to design and implement a solution where the hospital wants an automated process to upload any new medical images to Cloud Storage. On the basis of this statements which of the following statement is correct.

A. Create a Pub/Sub topic, and enable a Cloud Storage trigger for the Pub/Sub topic. Create an application that sends all medical images to the Pub/Sub topic.
B. Create a script that uses the gsutil command line interface to synchronize the on-premises storage with Cloud Storage. Schedule the script as a cron job.
C. In the Cloud Console, go to Cloud Storage. Upload the relevant images to the ap-propriate bucket.
D. Deploy a Dataflow job from the batch template, "Datastore to Cloud Storage" Schedule the batch job on the desired interval.

**Correct Answer: B**
**Section:**
**Explanation:**

Using sync for new images implies that you will continue to use your onprem and keep synchronizing it forever, Sync just once for the old images, new images go directly to google cloud via pub/sub, and eventually get rid of the onprem.

**QUESTION 116**
You are storing sensitive information in a Cloud Storage bucket. For legal reasons, you need to be able to record all requests that read any of the stored dat a. You want to make sure you comply with these requirements. What should you do?

A. Scan the bucket using the Data Loss Prevention API.
B. Enable Data Access audit logs for the Cloud Storage API.
C. Enable the Identity Aware Proxy API on the project.
D. Allow only a single Service Account access to read the data.

**Correct Answer: B**
**Section:**
**Explanation:**
Logged information
Your Google Cloud projects contain only the audit logs for resources that are directly within the Cloud project. Other Google Cloud resources, such as folders, organizations, and billing accounts, contain the audit logs for the entity itself.


Reference link- https://cloud.google.com/storage/docs/audit-logging

**QUESTION 117**
Your client has an on-premises data center. Due to technical limitations, they are unable to scale globally. They have decided to adopt the public cloud. However, they don't want to locked into any one vendor and, therefore, would like to work with multiple cloud providers. They have used open source container technologies and would like to continue using them.

A. Cloud Run which supports containers and can scale in a serverless fashion
B. Kubernetes that runs containers as their core workloads
C. AppEngine Flexible Environment which supports containers
D. Anthos that runs containers as their core workloads

**Correct Answer: D**
**Section:**
**Explanation:**
Anthos unifies the management of infrastructure and applications across on-premises, edge, and in multiple public clouds with a Google Cloud-backed control plane for consistent operation at scale.

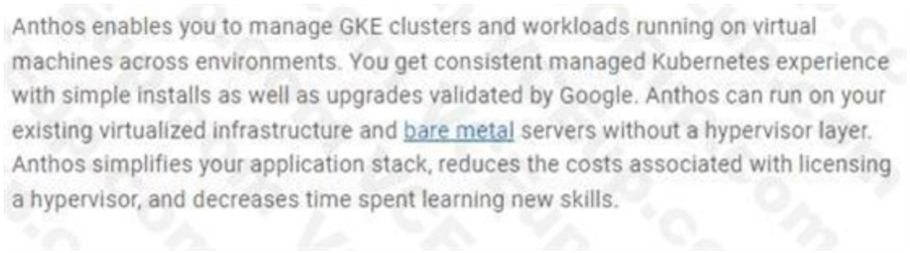


**QUESTION 118**
"With cloud messaging you can Customize and deliver messages accordingly to the predetermined time in the user's local time zone." Comment on the above statement.

A. This statement is undefined.
B. The above statement is partially true.
C. The above statement is completely false.
D. The above statement is completely true.

**Correct Answer: D**
**Section:**
**Explanation:**
Firebase Cloud Messaging:
Firebase Cloud Messaging (FCM) is a cross-platform messaging solution that lets you reliably send messages at no cost.
Using FCM, you can notify a client app that new email or other data is available to sync. You can send notification messages to drive user re-engagement and retention. For use cases such as instant messaging, a message can transfer a payload of up to 4000 bytes to a client app.
Key capabilities of Firebase Cloud Messaging:
Send notification messages or data messages: Send notification messages that are displayed to your user. Or send data messages and determine completely what happens in your application code.
Versatile message targeting: Distribute messages to your client app in any of 3 waysóto single devices, to groups of devices, or to devices subscribed to topics.
Send messages from client apps: Send acknowledgments, chats, and other messages from devices back to your server over FCM's reliable and battery-efficient connection channel.

**QUESTION 119**
You have deployed a new public web application that allows users to register and login with email ids, phone numbers, or user ids. You are seeing some unusual activity with user registrations and logins from a few IPs. A large number of accounts were created very quickly. Logins are also happening quickly thereafter from these new accounts. Different parts of the application are being explored, all of which are putting a heavy load on the application. What could be a problem and how can you solve it?

A. A hacker group has hired a bunch of people to create accounts and manually use the system. Use Cloud Asset Inventory to see if there have been changes in the inventory.
B. Bots are creating accounts and then using them. Use Google Cloud's Web App and API Protection (WAAP).
C. Bots are creating accounts and then using them. Use Identity-Aware Proxy to re-strict the users to known users.
D. Automated testing tools might still be running and creating accounts. Use Identity-Aware Proxy to restrict the users to known users.

**Correct Answer: B**
**Section:**
**Explanation:**
Bots attacking the application is the most likely scenario in this case. Using WAAP is the right protection plan: Anti-DDoS, anti-bot, WAF, and API protection help you protect against new and existing threats while helping you keep your apps and APIs compliant and continuously available.
https://cloud.google.com/solutions/web-app-and-api-protection

**QUESTION 120**
certain devices for cracks, rust, etc. Some of these issues are difficult to identify for a human and your company has seen increasing customer complaints - the customer has paid for an inspection and the field agent said there was no problem, but it later turned out there actually was. The team has come up with a proposal to engage AI to identify issues. On evaluating the existing system, it is seen that the mobile phone network connection is not

good or consistent. What solution can work for them?

A. Use AutoML Vision Edge models.
B. Use the Rust programming language instead of Python to identify issues like rust.
C. Use Cloud TPUs which will be able to do the analysis faster on the cloud. Thus re-sponses also will be fast.
D. Use TensorFlow to create custom models and deploy it as TensorFlow Lite mod-els.

**Correct Answer: A**
**Section:**
**Explanation:**
AutoML Vision Edge model can be deployed to one of several types of edge devices, such as mobile phones, ARM-based devices, and the Coral Edge TPU
https://cloud.google.com/vision/automl/docs/edge-quickstart

**QUESTION 121**
Your customer's IT team is in the process of modernizing their customer-facing applications. They've witnessed others getting good results from employing microservices, and they're keen to adopt it themselves. The first application that they are modernizing has about 5 different sub-parts, which they have identified will be the services. They also identify that each of them has different scale requirements - some services like user login are less frequently used while others like transac-tions are heavily used. What technical strategy would you recommend for them?

A. Containerize the services and orchestrate them with Google Kubernetes Engine.
B. Retain the original application in Compute Engine and scale it as needed using Managed Instance Groups.
C. Retain the original application as a backup and also for separately scaling the ser-vices, create new application binaries.
D. Retain the original application in Compute Engine and scale it as needed using Unmanaged Instance Groups.

**Correct Answer: A**
**Section:**
**Explanation:**
Containers and Kubernetes are ideal for the kind of requirement mentioned here - separate microservices that need to scale independently.
Google Kubernetes Engine (GKE) provides a managed environment for deploying, managing, and scaling your containerized applications using Google infrastructure. The GKE environment consists of multiple machines (specifically, Compute Engine instances) grouped together to form a cluster.
Reference link- https://cloud.google.com/kubernetes-engine/docs/concepts/kubernetes-engineoverview

**QUESTION 122**
Which of the following is/are core storage options available on the Google Cloud Platform?

A. Cloud Storage and Cloud Data Store
B. Cloud Spanner
C. Cloud SQL and Google Big Table
D. All of the above

**Correct Answer: D**
**Section:**
**Explanation:**
Google Cloud Platform has other storage options to meet your needs for structured, unstructured, transactional and relational data. Core storage options: Cloud Storage, Cloud SQL, Cloud Spanner, Cloud Data Store and Google Big Table.
Depending on your application, you might want to use one or several of these services to get the job done.

**QUESTION 123**
You are a program manager in a company you need to submit a bare metal solution order for a secure, high performance connection with a low-latency network fabric. What network information you need to submit the

order to Bare Metal
Solutions.

A. IP Ranges for example Client IP Address range used for communication between your Google Cloud and Bare Metal Solution environments.
B. Google Cloud Project Id that you are using with your bare metal solution environ-ment.
C. Total number of VLANs you need in your Bare Metal Solution Environment.
D. All of the above

**Correct Answer: D**
**Section:**
**Explanation:**
What Bare Metal Solution provides
Bare Metal Solution is a managed solution that provides purpose-built HPE or Atos bare-metal servers in regional extensions that are connected to Google Cloud by a managed, high-performance connection with a low-latency network fabric.
With Bare Metal Solution, Google Cloud provides and manages the core infrastructure, the network, the physical and network security, and hardware monitoring capabilities in an environment from which you can access all of the Google Cloud services. The core infrastructure includes secure, controlled-environment facilities, and power.
The Bare Metal Solution also includes the provisioning and maintenance of custom, sole-tenancy servers with local SAN, and smart hands support.
The network, which is managed by Google Cloud, includes a low-latency Partner Interconnect connection into the customer Bare Metal Solution environment.
The available Google Cloud services include private API access, management tools, support, and billing.

**QUESTION 124**
You are a program manager in a company and handling a project and you need to create a virtual machine on google cloud console that will be very simple to set up, by flipping a bit via command, API, or with developer console that gives you 30 seconds to shut down when you're preempted, allow you to save your work that also helps in the company budget upto 70-80% of less charges than the regular VMs.

A. Bare Metal Solutions
B. Preemptible Virtual Machines.
C. Google Cloud VM Instances
D. None of the above.

**Correct Answer: B**
**Section:**
**Explanation:**
Preemptible VMs have all these features
Simple configuration
Create a preemptible instance simply by flipping a bit via command, API, or developer console.
Easy extensibility
Attach GPUs and local SSDs to preemptible instances for additional performance and savings.
Graceful shutdown
Compute Engine gives you 30 seconds to shut down when you're preempted, letting you save your work in progress for later.
Large scale computing
Spin up as many instances as you need and turn them off when you're done. You only pay for what you use.
Quickly reclaim capacity
Managed instance groups automatically recreate your instances when they're preempted (if capacity is available).
Fixed pricing
Preemptible VMs have fixed pricing up to 80% off regular instances. They show up on your bill separately so you'll see just how much you're saving.

**QUESTION 125**
In terms of Infrastructure as a Service (IaaS) what are the benefits of it?

A. IaaS offers virtually infinite flexibility and scalability, enterprises can get their work done more efficiently, ensuring faster development life cycles.
B. IaaS resources are regularly available to businesses when they need them. As a result, enterprises reduce delays when expanding infrastructure and, alternatively, don't waste resources by overbuilding capacity.
C. IaaS resources are used on demand and enterprises only have to pay for the compute, storage, and networking resources that are actually used, IaaS costs are fairly predictable and can be easily contained and budgeted for.
D. All of the Above

**Correct Answer: D**
**Section:**
**Explanation:**
These are the feature of Infrastructure as a Service (IaaS)
It's economical
Because IaaS resources are used on demand and enterprises only have to pay for the compute, storage, and networking resources that are actually used, IaaS costs are fairly predictable and can be easily contained and budgeted for.
It's efficient
IaaS resources are regularly available to businesses when they need them. As a result, enterprises reduce delays when expanding infrastructure and, alternatively, don't waste resources by overbuilding capacity.
It boosts productivity
Because the cloud provider is responsible for setting up and maintaining the underlying physical infrastructure, enterprise IT departments save time and money and can redirect resources to more strategic activities.
It's reliable
IaaS has no single point of failure. Even if any one component of the hardware resources fails, the service will usually still remain available.
It's scalable
One of the biggest advantages of IaaS in cloud computing is the capability to scale the resources up and down rapidly according to the needs of the enterprise.
It drives faster time to market
Because IaaS offers virtually infinite flexibility and scalability, enterprises can get their work done more efficiently, ensuring faster development life cycles.

**QUESTION 126**
Your company has made plans to roll out OpenShift, a Kubernetes platform solution offered by IBM Red Hat, across all its on-premises and public cloud environments. Given that you are the lead architect responsible for your company's GCP deployments, what type of shared responsibility model will this deployment entail for you?

A. SaaS
B. On premises
C. PaaS
D. IaaS

**Correct Answer: D**
**Section:**
**Explanation:**
The key to remember here is that for a service provided (GCP in this case) to take responsibility for its PaaS, it must offer the service as a managed service. GCP offers its own Kubernetes platform called GKE. But OpenShift is not a Google-offered PaaS solution. As such, Google will not take responsibility for the back-end operations and design of your OpenShift environments. You will need to manage all the VMs that OpenShift will provision as part of its GCP deployment. So this is an IaaS deployment from a shared responsibility model perspective.

**QUESTION 127**
A client is currently running software on their on-premise systems that is bound by a certain type of license. They are allowed to run the software on virtualized machines. However, they cannot run them on virtualized machines that are shared by two different companies, teams, or projects. What option do they have on Google Cloud?

A. Google Cloud is a public cloud accessed by multiple customers.
B. Allocate a Bare Metal machine.
C. Setup exclusive login to the VM with self-generated security keys.
D. Allocate sole-tenant nodes

**Correct Answer: D**
**Section:**
**Explanation:**
Sole-tenancy lets you have exclusive access to a sole-tenant node, which is a physical Compute Engine server that is dedicated to hosting only your project's VMs. Use sole-tenant nodes to keep your VMs physically separated from VMs in other projects, or to group your VMs together on the same host hardware.
https://cloud.google.com/compute/docs/nodes/sole-tenant-nodes

**QUESTION 128**
Which of the following is / are true for Preemptible Instances.

A. Preemptible Instances have no Service Level Agreement (Compute Engine SLA).
B. Google Cloud Free Tier credits for compute engine do not apply to preemptible in-stances.
C. Preemptible instances can't live migrate to a regular VM instance, or be set to au-tomatically restart when there is a maintenance event.
D. All of the above.

**Correct Answer: D**
**Section:**
**Explanation:**
Preemptible instances function like normal instances but have the following limitations:
-> Compute Engine might stop preemptible instances at any time due to system events. The probability that Compute Engine will stop a preemptible instance for a system event is generally low, but might vary from day to day and from zone to zone depending on current conditions.
-> Compute Engine always stops preemptible instances after they run for 24 hours. Certain actions reset this 24-hour counter.
-> Preemptible instances are finite Compute Engine resources, so they might not always be available.
-> Preemptible instances can't live migrate to a regular VM instance, or be set to automatically restart when there is a maintenance event.
-> Due to the above limitations, preemptible instances are not covered by any Service Level Agreement (and, for clarity, are excluded from the Compute Engine SLA).
-> The Google Cloud Free Tier credits for Compute Engine do not apply to preemptible instances.

**QUESTION 129**
If you increase the size of a subnet in a custom VPC network, the IP addresses of virtual machines already on that subnet might be affected. Which options are Correct.

A. False
B. None of the above
C. True
D. Not Defined by Google Cloud Platform

**Correct Answer: A**
**Section:**
**Explanation:**
You can dynamically increase the size of a subnet in a custom network by expanding the range of IP addresses allocated to it. Doing that doesn't affect already configured VMs.

**QUESTION 130**
Which of the following statements is / are correct about Machine Learning?

A. Machine learning examples include chatbots and automated virtual assistants to automate routine customer service tasks and speed up issue resolution.
B. Machine learning automates the job of building statistical models with Human In-tervention.
C. Robotic process automation (RPA) can not be attached with ML.
D. None of the Above.

**Correct Answer: A**
**Section:**
**Explanation:**
Customer service
Machine learning examples include chatbots and automated virtual assistants to automate routine customer service tasks and speed up issue resolution.

**QUESTION 131**
You are a cloud architect in a software solution provider company, one of the client that is a Na-tional Bank who wants to build an application that deals with transactions processing, and it needs a relational database with petabyte of scale dat a. Which of the following Google Cloud Services will you use?

A. Cloud SQL
B. Cloud Bigtable
C. Cloud Spanner
D. Google Cloud BigQuery

**Correct Answer: C**
**Section:**
**Explanation:**
- Cloud Spanner is the online transaction processing solution that is relational and offers petabyte scalability. Cloud SQL is not designed for petabyte-scale data.

**QUESTION 132**
Your customer has a reporting tool that is only occasionally used by the leadership team. Usage of it is frequent - once a week, once a month, or once the quarter. They want to run this application in a cost-effective manner. What are the compute options available on Google Cloud which would be suitable? (Choose Two answer)

A. Cloud Run
B. Cloud App Engine Standard
C. Compute Engine
D. Kubernetes Engine

**Correct Answer: A, B**
**Section:**
**Explanation:**
Since the use of the tool is infrequent/intermittent, you can choose to compute options that are serverless. Both Cloud Run and Cloud App Engine Standard are serverless options that can shut down to zero. Since cost-effectiveness is a requirement, this will not cost anything during the periods it is not used.

**QUESTION 133**
You are working in a company where you need to store Terabytes of Image Data daily and process them e.g. Taking photos of the entire planet 24 hours every day with satellite and sending data to data centres to store and process it.
Which of the following would be the best combination for your infrastructure.
You are working in a company where you need to store Terabytes of Image Data daily and process them e.g. Taking photos of the entire planet 24 hours every day with satellite and sending data to data centres to store and process it.
Which of the following would be the best combination for your infrastructure.

A. Bare Metal Solutions with Google Cloud Storage.
B. Google Cloud Storage & Google Cloud Compute Engines
C. Google Cloud Storage & Preemptible VMs.
D. None of the Above

**Correct Answer: C**
**Section:**
**Explanation:**
The above is a real world example of a company named Planet, where they sent around 80+ satellites to take pictures of earth every day, 24 hours. They run around 40,000 preemptible VMs concurrently.
Preemptible instances function like normal instances but have the following limitations:
Compute Engine might stop preemptible instances at any time due to system events. The probability that Compute Engine will stop a preemptible instance for a system event is generally low, but might vary from day to day and from zone to zone depending on current conditions.
Compute Engine always stops preemptible instances after they run for 24 hours. Certain actions reset this 24-hour counter.
Preemptible instances are finite Compute Engine resources, so they might not always be available.
Preemptible instances can't live migrate to a regular VM instance, or be set to automatically restart when there is a maintenance event.
Due to the above limitations, preemptible instances are not covered by any Service Level Agreement (and, for clarity, are excluded from the Compute Engine SLA).
The Google Cloud Free Tier credits for Compute Engine do not apply to preemptible instances.


Reference link- https://cloud.google.com/compute/docs/instances/preemptible

**QUESTION 134**
DriveSuper Inc. teaches its clients to drive cars and bikes and helps them get their license. They are planning to build a mobile application where users can sign up, plan their schedules, and take stock of progress. They want the onboarding process to be smooth and frictionless, giving users a great experience from the get-go. They want this done as quickly as possible and not be expensive. What is their best option on Google Cloud?

A. Build the mobile app with Cloud SQL as the backend
B. Build the mobile app with Cloud Storage as the backend
C. Build the mobile application with Firebase as the backend
D. Build the mobile app with Cloud Spanner as the backend

**Correct Answer: C**
**Section:**
**Explanation:**
Firebase/Firestore is easy to build and is suitable for user information that could bvary in nature.

**QUESTION 135**
Which of the following is true while creating a boot persistent disk from a snapshot.

A. You cannot apply a snapshot to an existing persistent disk, or apply a snapshot to persistent disks that belong to a different project than that snapshot.
B. It is only possible to apply data from a snapshot when you first create a persistent disk.
C. After you create a snapshot of a boot persistent disk, you can apply data from that snapshot to new persistent disks.
D. All of the above.

**Correct Answer: D**
**Section:**
**Explanation:**
When you create a virtual machine (VM) instance, you must also create a boot disk for the VM. You can use a public image, a custom image, or a snapshot that was taken from another boot disk. When you create a boot disk, limit the disk size to 2 TB to account for the limitations of MBR partitioning.
Compute Engine automatically creates a boot persistent disk when you create an instance. If you require additional data storage space for your instances, add one or more secondary instance storage options.
You might need to create a standalone boot persistent disk and attach it to an instance later, or resize a boot persistent disk to improve performance and add more space for additional applications or operating system files.

That process is described in Add or resize a persistent disk.
As a best practice, do not use regional persistent disks for boot disks. In a failover situation, they do not force-attach to a VM.
After you create a snapshot of a boot persistent disk, you can apply data from that snapshot to new persistent disks. It is only possible to apply data from a snapshot when you first create a persistent disk. You cannot apply a snapshot to an existing persistent disk, or apply a snapshot to persistent disks that belong to a different project than that snapshot.

**QUESTION 136**
An application has become very popular and the number of requests/users is increasing quickly.
There is a meeting to figure out how to scale the systems so that they can accept user requests and still have the capacity to spare. What is the preferred option?

A. Circular Scaling takes a round-robin approach to allocate and destroy VMs.
B. Triangular Scaling takes an automated average of Cost, Effort, and Time.
C. Vertical Scaling
D. Horizontal Scaling

**Correct Answer: D**
**Section:**
**Explanation:**
Horizontal scaling, also called scaling out, adds new VMs to increase application capacity.

**QUESTION 137**
A customer in the European Union region is very clear that their data should not go outside the European Union. Their end users are spread all over the European U. They have to choose a storage option that serves all the users within Asia via web browsers as quickly as possible. Which storage option will work for them?

Multi-regions

| Multi-Region Name | Multi-Region Description |
|---|---|
| ASIA | Data centers in Asia |
| EU | Data centers within member states of the European Union* |
| US | Data centers in the United States |

A. Cloud Storage with a single region that is known to be within the European U
B. Cloud Filestore is connected to virtual machines which are guaranteed to be within the European UC . Cloud Storage with the multi-region option of European U
C. Cloud Storage with the dual-region option of European U

**Correct Answer: C**
**Section:**
**Explanation:**
Multi-region option will use multiple datacenters that are within the European Union. More regions will also help with lower latency since users are spread across the European U.
https://cloud.google.com/storage/docs/locations#considerations

**QUESTION 138**
In terms of Dockers and Kubernetes, which of the following statements are correct?

A. Kubernetes uses Docker to deploy, manage, and scale containerized applications.
B. Difference between Docker and Kubernetes relates to the role each play in con-tainerizing and running your applications
C. Kubernetes can be used with or without Docker.
D. All of the above.

**Correct Answer: D**
**Section:**
**Explanation:**
Kubernetes vs. Docker
Often misunderstood as a choice between one or the other, Kubernetes and Docker are different yet complementary technologies for running containerized applications.
Docker lets you put everything you need to run your application into a box that can be stored and opened when and where it is required. Once you start boxing up your applications, you need a way to manage them; and that's what
Kubernetes does.
Kubernetes is a Greek word meaning 'captain' in English. Like the captain is responsible for the safe journey of the ship in the seas, Kubernetes is responsible for carrying and delivering those boxes safely to locations where they can be used.
- Kubernetes can be used with or without Docker.
- Docker is not an alternative to Kubernetes, so it's less of a "Kubernetes vs. Docker" question. It's about using Kubernetes with Docker to containerize your applications and run them at scale.
- The difference between Docker and Kubernetes relates to the role each play in containerizing and running your applications.
- Docker is an open industry standard for packaging and distributing applications in containers.
- Kubernetes uses Docker to deploy, manage, and scale containerized applications.

**QUESTION 139**
In Google Cloud IAM: if a policy applied at the project level gives you Owner permissions, your access to an individual resource in that project might be restricted to View permission if someone applies a more restrictive policy directly to that resource. What is correct below the options

A. False
B. None of the above.
C. True
D. Not defined by GCP.

**Correct Answer: A**
**Section:**
**Explanation:**
Policies are a union of those applied to resources themselves and those inherited from higher levels in the hierarchy. If a parent policy is less restrictive, it overrides a more restrictive policy applied to the resource. If a parent policy is more restrictive, it does not override a less restrictive policy applied to the resource. Therefore, access granted at a higher level in the hierarchy cannot be taken away by policies applied at a lower level in the hierarchy.

**QUESTION 140**
Your customer is moving to Google Cloud. They have many teams, each working on many projects.
How should they organize resources?

A. Let each team have one shared Folder with multiple Projects within it so that there is a separation of concerns.
B. Let each Project have one Folder so that there is a clear separation of concerns.
C. Let each team have an Organization so that they can entirely manage themselves with their own identity.
D. Let each team have one shared Project so that it is easy to manage.

**Correct Answer: A**
**Section:**
**Explanation:**
The recommended approach is to have folders corresponding to teams/departments and they manage the projects within that.
-> Sharing a single project will cause a conflict of resources, billing, concerns, etc.
-> One folder per project is unnecessary overuse of abstraction/grouping.
-> Teams and projects in a company should ideally be centrally managed in a single Organization.

**QUESTION 141**
Considering Google Cloud Storage different Options which of the following is / are correct on the basis of their real world use cases?

A. Cloud Storage : Images, Large Media, files , backups.
B. Google Cloud BigTable : AdTech, Financial and IoT Data.
C. Cloud SQL : User Credentials, customer orders.
D. All of the Above.

**Correct Answer: D**
**Section:**
**Explanation:**
Cloud Datastore is the best for semi-structured application data that is used in app engines' applications. Bigtable is best for analytical data with heavy read/write events like AdTech, Financial or IoT data. Cloud Storage is best for structured and unstructured, binary or object data like images, large media files and backups. SQL is best for web frameworks and in existing applications like storing user credentials and customer orders. Cloud Spanner is best for large scale database applications that are larger than two terabytes; for example, for financial trading and e-commerce use cases. As I mentioned at the beginning of the module, depending on your application, you might use one or several of these services to get the job done.

**QUESTION 142**
The government has mandated that companies in a particular section of healthcare must retain all the data they collect for a period of 10 years in case an audit needs to be done. Your client, who is in that industry, needs to follow regulations. In addition, your client wants to do an analysis of the data quite frequently in the first year. They also don't want to be liable for any data beyond year 10. What would recommend for your customer?

A. Use Cloud Storage with nearline storage in year one and Coldline storage thereaf-ter. Use Object lifecycle management to move between storage types and delete them after 10 years.
B. Use Cloud Storage with standard storage in year one and Coldline storage there-after. Set a Cloud Scheduler trigger for 1 year to change storage types and 10 years to delete the data.
C. Use Cloud Storage with standard storage in year one and archival storage thereaf-ter. Use Object lifecycle management to move between storage types and delete them after 10 years.
D. Use Cloud Storage with standard storage in year one and Coldline storage there-after. Set a Cloud Tasks to trigger for 1 year to change storage types and 10 years to delete the data.

**Correct Answer: C**
**Section:**
**Explanation:**
Cloud storage supports Object Lifecycle Management. To support common use cases like setting a Time to Live (TTL) for objects, retaining noncurrent versions of objects, or "downgrading" storage classes of objects to help manage costs, Cloud Storage offers the Object Lifecycle Management feature.
Standard storage is recommended for frequently accessed data and Archive for data accessed less than once a year.
Nearline, Coldline, and Archive offer ultra-low-cost, highly-durable, highly available archival storage.
For data accessed less than once a year, Archive is a cost-effective storage option for the long-term preservation of data. Coldline is also ideal for cold storageódata your business expects to touch less than once a quarter. For warmer storage, choose Nearline: data you expect to access less than once a month, but possibly multiple times throughout the year.

**QUESTION 143**
You are discussing scaling requirements with a gaming company. When the game launches, they are expecting incoming data surges of 2 million users or more during weekends and holidays. Their onpremise systems have had issues scaling and they want your advice on solving the issue. What do you recommend?

A. Either Compute Engine VMs or Kubernetes nodes work, but it is better to keep a buffer of an extra 2 million users.
B. We can deploy a Pub/Sub to ingest data which will grow to absorb demand and pass it on to other stages.
C. We will allocate Compute Engine VMs estimating 80% capacity of 2 million users.
D. We will allocate Kubernetes nodes estimating 80% capacity of 2 million users.

**Correct Answer: B**
**Section:**
**Explanation:**
When there are huge surges in demand, it is preferable to use serverless technologies that automatically scale on demand. In this case, the key concern is data ingestion. Pub/Sub is a serverless system that can expand to

absorb such demand.

**QUESTION 144**
What type of cloud computing service provides raw compute, storage, and network, organized in ways that are familiar to physical data centers?

A. Database as a Service.
B. Platform as a Service.
C. Infrastructure as a Service.
D. Software as a Service.

**Correct Answer: C**
**Section:**
**Explanation:**
What is Infrastructure as a service :
IaaS (infrastructure as a service) is a computing model that offers resources on-demand to businesses and individuals via the cloud.
IaaS is attractive because acquiring computing resources to run applications or store data the traditional way requires time and capital. Enterprises must purchase equipment through procurement processes that can take months. They must invest in physical spaces: typically specialized rooms with power and cooling. And after deploying the systems, enterprises need IT, professionals, to manage them.
All this is challenging to scale when demand spikes or the business grows. Enterprises risk running out of capacity or overbuilding and ending up with infrastructure that suffers from low utilization.
These challenges are why IaaS use is steadily growing. Learn more about Compute Engine, Cloud Storage, etc.

**QUESTION 145**
You are looking for a one stop reference page for GCP support. What Page would you se-lect?

A. Compliance Hub
B. Google Cloud Platform Status
C. Support Hub
D. Pricing Page

**Correct Answer: C**
**Section:**
**Explanation:**
Google provides a page that brings together everything needed around support. Its called the Support Hub Reference link- https://cloud.google.com/support-hub

**QUESTION 146**
What load balancer type is supported with Cloud Armor security policies?

A. SSL Proxy, HTTP(S) and SSL
B. HTTP(S) and SSL
C. Regional SSL
D. HTTP(S) Only

**Correct Answer: D**
**Section:**
**Explanation:**
Google Cloud Armor security policies protect your application by providing Layer 7 filtering and by scrubbing incoming requests for common web attacks or other Layer 7 attributes to potentially block traffic before it reaches your load balanced backend services or backend buckets. Each security policy is made up of a set of rules that filter traffic based on conditions such as an incoming request's IP address, IP range, region code, or request headers.
-> Google Cloud Armor security policies are available only for backend services behind an external HTTP(S) load balancer. The load balancer can be in Premium Tier or Standard Tier.

-> Google Cloud Armor security policies and IP DENY lists and ALLOW lists are available only for HTTP(S) load balancing.
Reference link- https://cloud.google.com/armor/docs/security-policy-overview

**QUESTION 147**
Compute Engine provides machine type recommendations to help you optimize the re-source utilization of your virtual machine (VM) instances. What is this capability called?

A. App Engine
B. None of the above
C. Rightsizing Recommendations
D. Recommendation Engine

**Correct Answer: C**
**Section:**
**Explanation:**
Compute Engine provides machine type recommendations to help you optimize the resource utilization of your virtual machine (VM) instances. These recommendations are generated automatically based on system metrics gathered by the Cloud Monitoring service over the previous 8 days. Use these recommendations to resize your instance's machine type to use the instance's resources more efficiently. This feature is also known as rightsizing recommendations
Reference link- https://cloud.google.com/compute/docs/instances/apply-machine-typerecommendations- for-instances

**QUESTION 148**
All Google Cloud Platform services are associated with a project that is used to provide what functions?

A. Manage Container Deployments
B. Enable Services and APIs
C. Manage DNS Services
D. None of the Above

**Correct Answer: B**
**Section:**
**Explanation:**
Enable Services and APIs

A project organizes all your Google Cloud resources. **A project consists of a set of users; a set of APIs; and billing, authentication, and monitoring settings for those APIs.** So, for example, all of your Cloud Storage buckets and objects, along with user permissions for accessing them, reside in a project. **You can have one project, or you can create multiple projects and use them to organize your Google Cloud resources, including your Cloud Storage data, into logical groups.**

Reference link- https://cloud.google.com/storage/docs/projects

**QUESTION 149**
Cloud Data Loss Prevention (DLP) is a fully managed service designed to help discover, classify, and protect the most sensitive dat a. DLP provides three key features (Select Three Answers)

A. Classification
B. De-identification
C. De-classification
D. Inspection

E. Reinspection

**Correct Answer: A, B, D**
**Section:**
**Explanation:**
Classification. De-classification and Inspection
Classification is the process to inspect the data and know what data we have, how sensitive it is, and the likelihood. Inspection and classification happen here.
De-identification is the process of removing, masking, replacing information from data.
Reference link- https://cloud.google.com/dlp/docs

**QUESTION 150**
A developer in your IT team is cheating a bucket on Cloud Storage. He is receiving an error that the bucket name already exists. He has checked his project and the few other pro-jects in the organization, The name seems to be entirely unique, What would be the is-sue?

A. Bucket names ignore any "." in the name. Look for similar bucket names that have a "." in it.
B. Previously deleted bucket names in the same project cannot be reused. There must have been an older bucket with the same name.
C. Bucket names in Cloud storage have to be globally unique
D. Bucket name are case insensitive- look for bucket name in your org that have a different capitalization.

**Correct Answer: C**
**Section:**
**Explanation:**
Bucket names have to be unique across Google Cloud Platform [GCP], Including other organizations and projects.

**QUESTION 151**
Your company has signed up with a cloud provider and you will be using storage and vir-tual machines with the provider. The provider has provided your organization some expec-tations for what the service should perform at. What type of agreement provides a guar-antee of a certain level of service such as "Uptime"?

A. Performance Agreement
B. Interconnection Agreement
C. Warranty
D. Service Level Agreement

**Correct Answer: D**
**Section:**
**Explanation:**
Service Level Agreement (SLA)
A service level agreement (SLA) is a contract between a service provider (either internal or external) and the end user that defines the level of service expected from the service provider. Some common SLA's are uptime, Response Time, etc.

**QUESTION 152**
Which of the following are the current options for paid support in GCP? (Select Three An-swer)

A. Premier
B. Standard
C. Enhanced
D. Role
E. Premium

**Correct Answer: B, C, E**
**Section:**
**Explanation:**
Because GCP provides three options for paid support which are Standard, Enhanced and Premium.
Basic Support is included with your Google Cloud subscription which cover only Case, phone, and chat support for billing issues only
Reference link- https://cloud.google.com/support

**QUESTION 153**
What cloud service model would you want to select if you want to solve a particular busi-ness problem by providing CRM services in the cloud to your enterprises?

A. CaaS
B. SaaS
C. PaaS
D. IaaS

**Correct Answer: B**
**Section:**
**Explanation:**
SaaS ñ Software as a Service (SaaS) provides you a complete product that is run and managed by the service provider. You worry only about using the software and not about infrastructure.
SaaS provides the lowest level of flexibility and management control over the infrastructure.
(Example: Google Gsuite and MS O365)

**QUESTION 154**
App Engine has been deployed in your customers GCP cloud deployment. The customer would like to know more about the benefits of App Engine Flexible. Please advise them on the benefits of App Engine Flexible (Select Two Answers)

A. Supports autoscaling
B. Supports Docker containers
C. Supports mainframe connectivity
D. Source code is written in specific versions of the supported programming lan-guages only

**Correct Answer: A, B**
**Section:**
**Explanation:**
Autoscaling is supported in both Flexible and Standard environments. Flexible Environment does run a Docker container that includes a custom runtime or source code written in other programming languages.
Reference link - https://cloud.google.com/appengine/docs/the-appengine-environments

**QUESTION 155**
A startup client of yours does offline data processing for a few of its clients. They are mi-grating their applications and the associated data to Google Cloud. They have 100TB of data to move. They presently have a very small private data center setup connected to a local internet provider. The maximum bandwidth they are able to get is 100Mbps. How long will it take them to transfer the data over the internet if the transfer goes smoothly?

A. bout 12 days.
B. About 2 years.
C. About 100 days.
D. About 24 hours.

**Correct Answer: C**
**Section:**

**Explanation:**
The key reason I included this question is to clarify some terminologies that will be important for your estimates. The data size mentioned is a TB terabyte. Note the "byte". The speed is mentioned in Mbps, which is Megabits per second.
Note the "bits". 8 bits make a byte. So, to get the actual number of bits transferred, you need to multiply the TB number by 8.
Total data transferred (in bits) = 100 * 1,000,000,000,000 * 8 bits
Speed = 100Mbps = 100 * 1,000,000. i.e. 100 million bits are transferred per second.
Hence time taken to transfer all the data = Total Data/Speed = 8,000,000 seconds.
Number of seconds in a day = 24*60*60 = 86,400
Total time taken in days = 8,000,000/86,400 = 92.59 days
Reference link- https://cloud.google.com/architecture/migration-to-google-cloud-transferring-yourlarge-datasets#online_versus_offline_transfer

**QUESTION 156**
A small scale retailer has been collecting its point of sale transaction in a PostgreSQL Da-tabase. They have raised funding for a strategic expansion goal in the next year that will see them grow significantly in Asia, Europe, North America,
Which Database option should they choose in Google Cloud?

A. BigQuary
B. Spanner
C. Cloud SQL
D. Bigtable

**Correct Answer: B**
**Section:**
**Explanation:**
Spanner is a global scale Database that Support SQL querying, Similar to PostgreSQL, Which will be regional. So that will be a fairly smooth move, Since they have the time and the funding, they can plan for this migration.

**QUESTION 157**
A customer deploys an application to App Engine and needs to check for Open Web Appli-cation Security Project (OWASP) vulnerabilities. Which service should be used to accom-plish this?

A. Cloud Armor
B. Cloud Security Scanner
C. Binary Authorization
D. Forseti Security

**Correct Answer: B**
**Section:**
**Explanation:**
Web Security Scanner identifies security vulnerabilities in your App Engine, Google Kubernetes Engine (GKE), and Compute Engine web applications. It crawls your application, following all links within the scope of your starting URLs, and attempts to exercise as many user inputs and event handlers as possible.
Currently, Web Security Scanner only supports public URLs and IPs that aren't behind a firewall. Web Security Scanner currently supports the App Engine standard environment and App Engine flexible environments, Compute Engine instances, and GKE resources.
Reference link- https://cloud.google.com/security-command-center/docs/concepts-web-securityscanner-overview

**QUESTION 158**
Google Cloud Platform (GCP) provides three main compliance resource webpages. What are they?
(Select Three Answer)

A. Compliance Reports Manager

B. Support Hub
C. Compliance Offerings
D. GDPR Home Page
E. TechCentral

**Correct Answer: A, C, D**
**Section:**
**Explanation:**
Compliance Reports Manager, GDPR Home Page, Compliance Offerings
GCP provides three main compliance resource webpages
Compliance Reports Manager ñ https://cloud.google.com/security/compliance/compliance-reports-manager


Compliance Offerings ñ https://cloud.google.com/security/compliance/offerings

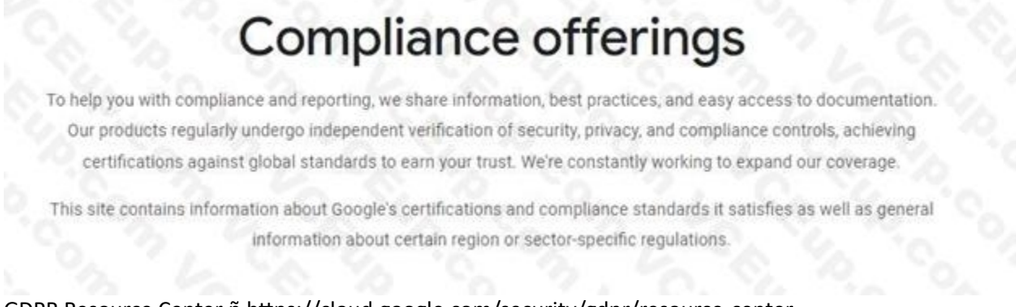


GDPR Resource Center ñ https://cloud.google.com/security/gdpr/resource-center
At Google Cloud, we champion initiatives that prioritize and improve the security and privacy ofcustomer personal data, and want you, as a Google Cloud customer, to feel confident using ourservices in light of GDPR requirements. If you partner with Google Cloud, we will support your GDPRcompliance efforts

**QUESTION 159**
What cloud deployment model is generally deployed between organizations such as non-profits, hospitals or even enterprises that share similar requirements or interests?

A. Hybrid
B. Community
C. Private
D. Public

**Correct Answer: B**
**Section:**
**Explanation:**
Community Cloud ñ The cloud infrastructure is planned for selective use by a particular community of consumers from organizations that have mutual interests like security needs, policy, and compliance considerations.
Reference link- https://csrc.nist.gov/glossary/term/community_cloud

**QUESTION 160**
What service is a fully managed real-time messaging service that allows you to send and receive messages between independent applications.

A. Cloud Datastore
B. Cloud Pub/Sub
C. Cloud DNS
D. Cloud BigTable
E. Cloud Spanner

**Correct Answer: B**
**Section:**
**Explanation:**
Google Cloud Pub/Sub is a scalable, durable event ingestion and delivery system.
-> Pub/Sub allows services to communicate asynchronously, with latencies on the order of 100 milliseconds.
-> Pub/Sub is used for streaming analytics and data integration pipelines to ingest and distribute data. It is equally effective as messaging-oriented middleware for service integration or as a queue to parallelize tasks.
-> Pub/Sub enables you to create systems of event producers and consumers, called publishers and subscribers. Publishers communicate with subscribers asynchronously by broadcasting events, rather than by synchronous remote procedure calls (RPCs).
Reference link- https://cloud.google.com/pubsub/docs/overview

**QUESTION 161**
A customer is migrating there on-promises data analytics solution to Google Cloud. The current solution has a lot of data being read form and written to disk. The performance of this approach has occasionally been a bottleneck for a scale of operations that your cus-tomer has. The application is fault tolerant and can with stand machine going down fre-quently. In moving to Google Cloud they are asking your advice on any way to improve performance?

A. Use Big Query Which has very fast data access and analysis
B. Use Cloud Storage which can be central, scalable storage
C. Use local SSDs with the VMs
D. Use Persistent Disk with the VMs

**Correct Answer: C**
**Section:**
**Explanation:**
Local SSDs are attached to the VM and have very high throughput. However, when the VM shuts down, The local SSD is also shut down, Since our Workload here is foult tolerant, than is not an issue.

**QUESTION 162**
Virtual Machine vCPU and memory usage for each of these categories can receive one of the following discounts? (Select Three Answer)

A. Military Discounts
B. Spot Instances
C. Committed-Use
D. Sustained-Use

E. Preemptible VMs

**Correct Answer: C, D, E**
**Section:**
**Explanation:**
Sustained, Committed and Preemptible vCPU and memory usage for each of these categories can receive discounts VM vCPU and memory usage for each of these categories can receive discounts Sustained-use discountsóGoogle offers up to 30% off for workloads that run for most of the billing month on GCP services.
Committed-use discountsóusers can save up to 57% by committing to use an instance for a certain time period, with no upfront payment and with the flexibility to change instances during the commitment period.
Preemptible VMsósimilar to the concept of AWS spot instances, Google offers up to 79% off for Virtual Machines that may be shut down at any time and replaced by others.
Reference link- https://cloud.google.com/compute/docs/sustained-use-discounts
Reference linkñ https://cloud.google.com/compute/docs/instances/signing-up-committed-usediscounts
Reference linkñ https://cloud.google.com/compute/docs/instances/preemptible

**QUESTION 163**
Your customer is moving from AWS to Google Cloud. Data also needs to be moved. There is about 50TB of dat a. On AWS, the data resides in an S3 bucket. It is going to be moved to Cloud Storage. Data is also being continuously generated on S3 prior to the cutover. It is preferable that this is also periodically transferred. What is the best way to move the data?

A. Use the gsutil command-line option
B. Use the Google Cloud console to drag and drop the files easily
C. Use the Storage Transfer Service
D. Use a Transfer Appliance

**Correct Answer: C**
**Section:**
**Explanation:**
Storage Transfer Service provides options that make data transfers and synchronization easier. We can also schedule one-time transfer operations or recurring transfer operations.


Reference link- https://cloud.google.com/storage-transfer/docs/overview
Reference link- https://cloud.google.com/architecture/transferring-data-from-amazon-s3-to-cloudstorage-using-vpc-service-controls-and-storage-transfer-service

**QUESTION 164**
Which of the following statements describe the features of a preemptible VM in-stance? (Select Three Answer)

A. Instance is alive for no more than 12 hours
B. Can be pre-empted with a 30 minute notice
C. Can be pre-empted with a 30 second notice
D. Discounted Significantly
E. Instance is alive for no more than 24 hours
F. Can use free tier credits

**Correct Answer: C, D, E**
**Section:**
**Explanation:**
Instance is alive for no more than 24 hours, Can be pre-empted with a 30 second notice, Discounted Significantly.
Preemptible VM is an instance that you can create and run at a lower cost than normal instances.
However, Compute Engine might stop (pre-empt) these instances if it requires access to those resources for other tasks. Preemptible instances are excess Compute Engine capacity, so their availability varies with usage.
Live at most 24 hours Can be pre-empted with a 30 second notification via API and are Discounted significantly Reference link- https://cloud.google.com/compute/docs/instances/preemptible

**QUESTION 165**
Your Coustomer's Organization has decided to move to the cloud. They currently run VMs onpromise but their goal on Google cloud is to run containers, primarily on Google Kuber-nete's Engine.
They have a lease for their private data center for another year that they have already paid for. What could be strategy they could adopt in migrating?

A. Jump and Ramp.
B. Improve and Move.
C. Rip and Replace.
D. Left and Shift.

**Correct Answer: B**
**Section:**
**Explanation:**
Since they have already paid for data center for another year. They have the time and resources to work with, They can make the change to their workloads locally/on-promise Improve and Migrate Move to Google Cloud later on.

**QUESTION 166**
A Customer has their current SAP systems using Microsoft SQL Server as the Database. They are migrating to Google Cloud and also preparing to later migrate to the latest ver-sion of SAP. The entire IT team is being directed to focus on the migration to the new ver-sion of SAP. The new version of SAP does not use Microsoft SQL Server as the Database, Any but the most critical IT management tasks are being deprioritized, How should they migrate their current database to Google Cloud?

A. Spanner
B. Bare Metal
C. BigQuery
D. Cloud SQL

**Correct Answer: D**
**Section:**
**Explanation:**
Cloud SQL supports SQL Server, Since the IT team's attention is being focused on other activities, they will have less time for existing admin tasks, It would be best to take a managed/hosted version.

**QUESTION 167**
An organization operates their entire IT infrastructure from Google Cloud.

What should they do to prepare for data breaches?

A. Reduce reliance on multi-factor authentication
B. Data security is Google's responsibility, so preparation is minimal
C. Create an incident plan to mitigate impacts
D. Strengthen their data center perimeter security

**Correct Answer: C**
**Section:**

**QUESTION 168**
An organization wants to build an entirely new infrastructure and applications in the cloud.
Which application modernization approach should the organization use?

A. Move the application to the cloud, and then change it.
B. Change their application, and then move it to the cloud.
C. Invent in greenfield.
D. Invent in brownfield.

**Correct Answer: C**
**Section:**
**Explanation:**
A Greenfield approach is a brand-new implementation , where companies then add their needed configurations and customizations. This approach provides a clean slate to start from, does not carry over needless customizations and technical debt, and provides a solid foundation for business process re-engineering.
A greenfield deployment is the design, installation and configuration of computer infrastructure where none existed before, for example, in a new office. In contrast, a brownfield deployment is an upgrade or addition to existing infrastructure using legacy components.

**QUESTION 169**
An organization wants to use Apigee to manage all their application programming interfaces (APIs).
What will Apigee enable the organization to do?

A. Increase application privacy
B. Measure and track API performance Most Voted
C. Analyze application development speed
D. Market and sell APIs

**Correct Answer: B**
**Section:**
**Explanation:**
Apigee's API Monitoring enables you to track your APIs to make sure they are up and running correctly. API Monitoring provides near real-time insights into API traffic and performance, to help you quickly diagnose and solve issues as they arise.
Apigee works with APIs not necessarily applications. It allows organizations to gain actionable insights across the entire API value chain and monetize API products and maximize the business value of digital assets.
https://cloud.google.com/apigee#section-11

**QUESTION 170**
A cloud-native organization is not meeting their service level objective (SLO) but has not exhausted their error budget.
What should the organization prioritize?

A. Innovation to improve user experience
B. Hardware reliability to improve availability
C. Stability to avoid prolonged user downtime
D. Speed to release new features

**Correct Answer: C**
**Section:**
**Explanation:**
Both Devs and SRE team must ensure that the error budget does not become exhausted. To avoid it, releases have to stop for the time being until the error budget resets. The team would have to reprioritize to focus on reliability to get it back to an acceptable state.

**QUESTION 171**
An organization finds that the amount of cash in their vending machines doesn't match the value of items sold. They have decided to upgrade their vending machines with cloud-based mobile payment systems.
How could the organization benefit from this upgrade?

A. They could relax data access permissions.
B. They could reduce their error budget overspend.
C. They could improve their perimeter security.
D. They could view data history to see transactions.

**Correct Answer: D**
**Section:**

**QUESTION 172**
An organization's developers are growing increasingly frustrated by the limitations of their onpremises infrastructure.
How would they benefit from leveraging cloud technology?

A. They can expect 100% service availability.
B. They can avoid the limitations of serverless computing.
C. They can have new tools to innovate and optimize resource usage.
D. They can optimize maintenance for their on-premises infrastructure.

**Correct Answer: C**
**Section:**
**Explanation:**
Google cloud have vast majority of products/tools that you can use to innovate. Additionally, there are products in google that scale automatically based from usage (Ex. App Engine, Cloud Run, etc.)

**QUESTION 173**
An online retail organization wants to optimize their service.
What is an example of unstructured data that they can use to make decisions?

A. Customer survey comments
B. Seller location coordinates
C. Product sales trends
D. Warehouse inventory records

**Correct Answer: A**

**Section:**
**Explanation:**
https://cloud.google.com/storage/docs/requester-pays

**QUESTION 174**
An organization wants to add a new function to their application. They want to write the code and let the public cloud provider handle the infrastructure.
Which infrastructure solution should they use?

A. Virtual machines
B. Bare Metal Solution
C. Serverless computing
D. Container Registry

**Correct Answer: C**
**Section:**
**Explanation:**
Serverless computing , as public cloud prouder(eg. google) will mange the infra things

**QUESTION 175**
Several departments in an organization are working together on a project. The organization wants to customize access to resources for each department.
What is the quickest and most efficient way to achieve this?

A. By mapping IAM roles to job functions for each department
B. By assigning IAM primitive roles to each employee
C. By applying least-privilege to roles for each employee
D. By creating a single shared service account for all departments

**Correct Answer: A**
**Section:**

**QUESTION 176**
An organization wants to create a new application in the cloud to replace an existing on-premises application.
Which application modernization approach should the organization use?

A. Move the application to the cloud, and then change it.
B. Change their application, and then move it to the cloud.
C. Invent in greenfield.
D. Invent in brownfield.

**Correct Answer: D**
**Section:**
**Explanation:**
This approach carries over as much custom components as possible from the source system and minimizes initial reengineering efforts.

**QUESTION 177**
A manager wants to review Google Cloud data access among their employees.
Who is responsible for defining data access policies?

A. Cloud Identity
B. Google Cloud Customer Care team
C. Their organization's IT team
D. Their organization's end users

**Correct Answer: C**
**Section:**
**Explanation:**
Cloud Identity and Access Management (IAM) helps customers to define fine-grained access policies and precisely control access to Google Cloud-hosted data.

**QUESTION 178**
An organization wants full control of their virtual machine infrastructure for a custom home-grown application with a product that autoscales and automatically updates.
Which Google Cloud product or solution should the organization use?

A. Cloud Build
B. Cloud Run
C. Compute Engine
D. App Engine

**Correct Answer: C**
**Section:**
**Explanation:**
Compute Engine will allow you to have full control of their VM infrastructure and you can autoscale and also apply automatic updates.

**QUESTION 179**
An organization wants to upskill their IT staff.
How can they do this in a transformational way?

A. Prioritize training current employees instead of hiring new recruits with cloud experience.
B. Prioritize giving privileged access to third-party partners and contractors to fill IT knowledge gaps.
C. Create a culture of self-motivated, isolated learning with official training materials.
D. Create a culture of continuous peer-to-peer learning with official training materials.

**Correct Answer: D**
**Section:**

**QUESTION 180**
How is privacy defined in the context of cloud technology?

A. Restrictions on data access and sharing
B. Procedures to authenticate user identity
C. Susceptibility to data breaches and cyber attacks
D. Compliance with regulatory standards

**Correct Answer: A**
**Section:**

**QUESTION 181**
How does a large hotel chain benefit from storing their customer reservation data in the cloud?

A. On-premises hardware access to transaction data
B. Real-time data transformation at scale within an on-premises database
C. Real-time business transaction accuracy at scale
D. Physical hardware access during peak demand

**Correct Answer: C**
**Section:**

**QUESTION 182**
An organization needs a platform to create custom end-to-end artificial intelligence models.
Which Google Cloud product or service should the organization use?

A. Dataproc
B. Compute Engine
C. Recommendations AI
D. Vertex AI

**Correct Answer: D**
**Section:**
**Explanation:**
Recommendations AI enables you to build an end-to-end personalized recommendation system based on state-of-the-art deep learning ML models, without a need for expertise in ML or recommendation systems. With Vertex AI, both AutoML training and custom training are available options. Whichever option you choose for training, you can save models, deploy models, and request predictions with Vertex AI.
https://cloud.google.com/vertex-ai

**QUESTION 183**
An organization wants to use BigQuery data analytics to understand their website performance, but wants to move only some data into the cloud.
Which environment should the organization use?

A. Private cloud
B. On-premises
C. Multi-cloud
D. Hybrid cloud

**Correct Answer: D**
**Section:**
**Explanation:**
The assumption should be made that there is still a private network involved. Hybrid clouds always include a private cloud and are typically managed as one entity. Multi-clouds always include more than one public cloud service, which often perform different functions.

**QUESTION 184**
Why do organizations often struggle to scale their on-premises application infrastructure?

A. Scaling compute instances could breach compliance and/or regulation
B. Increasing compute capacity is time-consuming and costly

C. Their serverless compute functions struggle to meet the demand
D. Their multi-cloud architecture is complex and expensive

**Correct Answer: B**
**Section:**

**QUESTION 185**
What does Cloud Debugger help an organization do?

A. Implement code updates in real time without affecting the service level objective (SLO).
B. Inspect source code in real time without affecting user downtime.
C. Manage code and accelerate application development.
D. Analyze live source code during user downtime.

**Correct Answer: B**
**Section:**
**Explanation:**
Cloud Debugger is a feature of Google Cloud Platform that lets you inspect the state of an application, at any code location, without stopping or slowing down the running app. Cloud Debugger makes it easier to view the application state without adding logging statements.

**QUESTION 186**
Why is data stored in Google Cloud secure and private?

A. Data is encrypted by the Security Command Center
B. Data is encrypted by Cloud Data Loss Prevention
C. Data is encrypted by default
D. Data is encrypted when an appropriate tag is applied

**Correct Answer: C**
**Section:**
**Explanation:**
https://cloud.google.com/docs/security/encryption/defaultencryption#:~:text=Google%20uses%20the%20Advanced%20Encryption,to%202015%20that%20use%20AES128

**QUESTION 187**
An organization needs to categorize text-based customer reviews on their website using a pretrained machine learning model.
Which Google Cloud product or service should the organization use?

A. Cloud Natural Language API
B. Dialogflow
C. Recommendations AI
D. TensorFlow

**Correct Answer: A**
**Section:**
**Explanation:**
https://cloud.google.com/natural-languageUse entity analysis to find and label fields within a documentóincluding emails, chat, and socialmediaóand then sentiment analysis to understand customer opinions to find actionable product andUX insights.

**QUESTION 188**
An organization is moving away from an on-premises infrastructure. Instead, they want to create, access, and share information virtually in the cloud.
What should the organization consider?

A. Built-in security when moving their data to the cloud
B. Replacing their perimeter security with data encryption keys
C. Optimizing cost-management with a capital expenditure model
D. Increased hardware capacity when moving their data to the cloud

**Correct Answer: A**
**Section:**

**QUESTION 189**
How does Google Cloud ensure that customer data remains secure and private when at rest?

A. By aggregating training data for customers within each industry
B. By automatically locking files containing suspicious code
C. By auditing platform privacy practices against industry standards
D. By providing privacy reviews for critical customer applications

**Correct Answer: C**
**Section:**
**Explanation:**
Google Cloud commitment to keep the data secure and private:
1. Org owns the data and not Google
2. Google does not sell data to 3rd parties
3. All customer data is encrypted by default
4. Google Cloud guards insider against your data
5. No backdoor access to any govt. entity
6. Google's privacy practices are audited against international standards

**QUESTION 190**
An organization delivers a proactive healthcare service. They want to efficiently and automatically collect patient data.
What should the organization encourage the patients to do?

A. Use at-home health screening devices and then upload their health data daily
B. Wear Internet of Things (IoT) devices that upload their health data in real time
C. Self-assess their health data and then document and upload it in real time
D. Visit a nurse who will use Internet of Things (IoT) devices to collect and upload their health data

**Correct Answer: B**
**Section:**

**QUESTION 191**
When an organization adopts cloud technology, how does their total cost of ownership (TCO) shift?

A. Away from cost management toward capital expenditure

B. Away from operational expenditure toward cost management
C. Away from capital expenditure toward operational expenditure
D. Away from operational expenditure toward capital expenditure

**Correct Answer: C**
**Section:**

**QUESTION 192**
An organization is altering their gaming product so that it is compatible with cloud technology.
What can they expect when moving from traditional technology to cloud technology?

A. No change to existing responsibilities
B. A shift toward OpEx
C. A shift toward using structured data
D. Increased hardware maintenance

**Correct Answer: B**
**Section:**

**QUESTION 193**
An organization wants a cost-effective relational database.
Which Google Cloud service should the organization use?

A. Cloud Storage
B. BigQuery
C. Cloud SQL
D. Dataflow

**Correct Answer: C**
**Section:**

**QUESTION 194**
How is service availability measured in the context of cloud technology?

A. Number of available regions
B. Percentage of uptime
C. Speed of response time
D. Number of downtime incidents

**Correct Answer: B**
**Section:**

**QUESTION 195**
An organization needs to run frequent updates for their business app.
Why should the organization use Google Kubernetes Engine (GKE)?

A. Customer expectations can be adjusted without using marketing tools

B. Seamless changes can be made without causing any application downtime.
C. GKE handles version control seamlessly and out of the box
D. GKE is well suited for all monolithic applications

**Correct Answer: B**
**Section:**
**Explanation:**
https://cloud.google.com/architecture/migrating-a-monolithic-app-to-microservices-gke

**QUESTION 196**
A food delivery service needs access to real-time menu information from all partner restaurants.
They also need to share customer order information with the restaurants in real time.
What should the organization use?

A. Site reliability engineering (SRE)
B. An application programming interface (API)
C. A customized machine learning model
D. A multi-regional database

**Correct Answer: B**
**Section:**

**QUESTION 197**
An organization wants to migrate legacy applications currently hosted in their data center to the cloud. The current architecture dictates that each application needs its own operating system (OS) instead of sharing an OS.
Which infrastructure solution should they choose?

A. Virtual machines
B. Open source
C. Serverless computing
D. Containers

**Correct Answer: A**
**Section:**
**Explanation:**
Virtual machines - you can install customized OS
Containers - about applications
Virtualization enables you to run multiple operating systems on the hardware of a single physical server, while containerization enables you to deploy multiple applications using the same operating system on a single virtual machine or server. Serverless computing would be no OS required and the open source operating system allows the use of code that is freely distributed and available to anyone and for commercial purposes such as Linux and Free BSD.

**QUESTION 198**
An organization is training a machine learning model to predict extreme weather events in their country.
How should they collect data to maximize prediction accuracy?

A. Collect all weather data evenly across all cities
B. Collect all weather data primarily from at-risk cities
C. Collect extreme weather data evenly across all cities

D. Collect extreme weather data primarily from at-risk cities

**Correct Answer: A**
**Section:**
**Explanation:**
Collect all weather data evenly across all cities. Mainly because it seems that the emphasis for data collection for ML is to make sure there are no holes in your data collection.

**QUESTION 199**
An organization wants to move from a tactical cloud adoption approach to a transformational approach.
How should they change their cloud security?

A. Provide staff identities using only Google Cloud authentication.
B. Provide multiple layers of network security using a zero-trust model.
C. Emphasize strong perimeter security and trust in their private network.
D. Emphasize three main Identity Access Management roles: owner, editor, and viewer.

**Correct Answer: B**
**Section:**

**QUESTION 200**
An organization wants to search hundreds of scanned documents for key information like dates, names, and other specific words.
Why should the organization use application programming interfaces (APIs)?

A. To replace the scanned documents with an online survey
B. To ingest data in real time and encrypt unmatched words
C. To create digital versions of the documents and locate key information
D. To transform the documents into unstructured data.

**Correct Answer: D**
**Section:**
**Explanation:**
The text from the PDF/scanned documents/images gets converted into JSON (unstructured file)which will be further used for search.

**QUESTION 201**
An organization wants to write and run small pieces of code in a serverless way that respond to events like huge discounts.
Which Google Cloud compute solution should the organization use?

A. Google Kubernetes Engine
B. Cloud Functions
C. Bare Metal Solution
D. Compute Engine

**Correct Answer: B**
**Section:**

**QUESTION 202**
What does Cloud Logging help an organization do?

A. Analyze live source code and log code updates.
B. Deploy infrastructure as code.
C. Analyze logs and accelerate application troubleshooting.
D. Manage storage of custom VM images.

**Correct Answer: C**
**Section:**

**QUESTION 203**
What is artificial intelligence?

A. Any system that ingests data in real time
B. Any system that automatically structures data
C. Any system capable of a task that requires smart analytics to generate predictions
D. Any system capable of a task that normally requires human cognition

**Correct Answer: D**
**Section:**

**QUESTION 204**
An organization has an on-premises IT infrastructure. Their customer-facing application repeatedly fails during peak usage.
What could be causing this issue?

A. A serverless compute function struggles to scale.
B. The application contains unclean data.
C. They don't have enough servers to meet the demand.
D. The application is only configurable on-premises.

**Correct Answer: C**
**Section:**

**QUESTION 205**
Which policy helps Google Cloud keep customer data private?

A. Google tests the service availability of customer applications.
B. Google does not use customer data for advertising purposes.
C. Google migrates customer data to an offline server when a threat is detected.
D. Google does not allow customers to change encryption keys.

**Correct Answer: B**
**Section:**

**QUESTION 206**
Why should an organization consider the total cost of ownership (TCO) when moving from onpremises to the cloud?

A. To evaluate error budget
B. To understand service level availability

C. To evaluate return on investment
D. To calculate required compute power

**Correct Answer: C**
**Section:**

**QUESTION 207**
What is monitoring within the context of cloud operations?

A. Observing cloud expenditure in real time to ensure that budgets are not exceeded
B. Collecting predefined and custom metrics from applications and infrastructure
C. Tracking user activities to guarantee compliance with privacy regulations
D. Tracing user location to document regional access and utilization

**Correct Answer: B**
**Section:**

**QUESTION 208**
An organization notices that some of their cloud expenditures are too high.
What should the organization do to control costs?

A. Streamline the hardware procurement process to reduce costs.
B. Share cost views with the departments to establish more accountability.
C. Change the cost model from operational expenditure to capital expenditure.
D. Ensure that all could resources are tagged with a single tag.

**Correct Answer: B**
**Section:**

**QUESTION 209**
An organization needs to migrate specialized workloads to the cloud while maintaining their existing complex licensing and architecture.
What Google Cloud solution should the organization use?

A. Compute Engine
B. Bare Metal Solution
C. Cloud Run
D. Cloud Functions

**Correct Answer: B**
**Section:**
**Explanation:**
"This solution provides a path to modernize your application infrastructure landscape, while maintaining your existing investments and architecture. With Bare Metal Solution, you can bring your specialized workloads to Google Cloud, allowing you access and integration with GCP services with minimal latency."

**QUESTION 210**
How would a global organization benefit from managing their data with Cloud Spanner?

A. Cloud Spanner is optimized for cold storage
B. Cloud Spanner replicates data across regions in real time
C. Cloud Spanner is optimized to ingest unstructured data
D. Cloud Spanner visualizes and analyzes data in real time

**Correct Answer: B**
**Section:**
**Explanation:**
Spanner is Google's scalable, multi-version, globally-distributed, and synchronously-replicated database.

**QUESTION 211**
How would an organization benefit from using Looker?

A. Optimal identity and access management
B. Leading serverless warehousing technology
C. Robust data roll-back accuracy
D. Advanced business intelligence and analytics

**Correct Answer: D**
**Section:**
**Explanation:**
Looker is a business intelligence software and big data analytics platform that helps you explore, analyze and share real-time business analytics easily.

**QUESTION 212**
An organization needs to store structured, semi-structured, and unstructured data in its raw, native format in the same repository.
Which cloud data management solution should the organization use?

A. Data field
B. Data lake
C. Database
D. Data warehouse

**Correct Answer: B**
**Section:**
**Explanation:**
A data lake can store all types of data with no fixed limitation on account size or file and with no specific purpose defined yet. The data comes from disparate sources and can be structured, semistructured, or even unstructured. Data-lake data can be queried as needed.
https://cloud.google.com/learn/what-is-a-data-lake
A data lake is a centralized repository designed to store, process, and secure large amounts of structured, semistructured, and unstructured data. It can store data in its native format and process any variety of it, ignoring size limits.

**QUESTION 213**
An organization wants to leverage tooling and automation as part of its new DevOps philosophy.
Which operational challenge will this resolve?

A. Repetitive manual tasks that hinder workflows
B. Time-consuming supervision of creative tasks
C. Distribution and supply-chain issues

D. Defective technical equipment that limits innovation

**Correct Answer: A**
**Section:**

**QUESTION 214**
An organization recently launched a virtual customer support agent, generating vast amounts of text and speech data.
Why should they use a cloud data warehouse to interpret this data?

A. To natively visualize both types of data using a dashboard in real time
B. To ingest and analyze structured and unstructured data at scale, in real time
C. To secure data transmission between cloud and on-premises environments
D. To transform data from structured to unstructured

**Correct Answer: B**
**Section:**
**Explanation:**
Real-time data ingestion and updates. A simple and universal solution for continually ingesting your enterprise data into popular cloud-based data warehouses in real time.
https://www.qlik.com/us/cloud-data-migration/cloud-data-warehouse

**QUESTION 215**
How does Cloud SQL help organizations create business insights?

A. Integrates with business intelligence and analytics platforms
B. Generates predictions using machine learning models
C. Generates real-time charts and intelligent analytics
D. Transforms business data from unstructured to structured

**Correct Answer: A**
**Section:**
**Explanation:**
https://cloud.google.com/sql/docs/postgres/using-query-insights

**QUESTION 216**
What is logging within the context of cloud technology?

A. Writing application and operating system events as text
B. Monitoring network and resource limitations
C. Tracking source code across an organization
D. Recording infrastructure and hardware expenditure

**Correct Answer: A**
**Section:**
**Explanation:**
Cloud Logging is a fully managed service that allows you to store, search, analyze, monitor, and alert on logging data and events from Google Cloud and Amazon Web Services

**QUESTION 217**
An organization needs to search an application's source code to identify a potential issue. The application is distributed across multiple containers.

Which Google Cloud product should the organization use?

A. Google Cloud Console
B. Cloud Trace
C. Cloud Monitoring
D. Cloud Logging

**Correct Answer: B**
**Section:**
**Explanation:**
Cloud Trace is supposed to be the correct answer. It's an application performance management tool.
It's a Google solution for monitoring application performance. It is a distributed tracing system that helps developers debug or fix and optimize their code

**QUESTION 218**
How can a streaming service meet global compliance requirements using the cloud?

A. By automatically encrypting personally identifiable information
B. By obtaining a business license to operate in a new market
C. By allowing users to disable two-factor authentication
D. By adhering only to data polices of the country in which the head office is registered

**Correct Answer: A**
**Section:**

**QUESTION 219**
An organization wants to use all available data to offer predictive suggestions on their website that improve over time.
Which method should the organization use?

A. Data automation
B. Trends analysis
C. Machine learning
D. Multiple regression

**Correct Answer: C**
**Section:**

**QUESTION 220**
An organization wants to transform multiple types of structured and unstructured data in the cloud from various sources. The data must be readily accessible for analysis and insights.
Which cloud data storage system should the organization use?

A. Relational database
B. Private data center
C. Data field
D. Data warehouse

**Correct Answer: D**
**Section:**

**Explanation:**
It supports real-time insights. A data warehouse is an enterprise system used for the analysis and reporting of structured and semi-structured data from multiple sources,
https://cloud.google.com/learn/what-is-a-data-warehouse

**QUESTION 221**
A retail company stores their product inventory in a legacy system. Often, customers find products on the company's website and want to purchase them in-store.
However, when they arrive, they discover that the products are out of stock.
How could the company benefit from using an application programming interface (API)?

A. To create personalized product recommendations for customers
B. To optimize their on-premises legacy system stability
C. By manually linking each inventory system to the website on a case-by-case basis
D. By programmatically connecting the inventory system to their website

**Correct Answer: D**
**Section:**
**Explanation:**
By programmatically connecting the inventory system to their website The issue is the website shows an item is available at the store, but when the customer gets to the store, they find out that item is out of stock.

**QUESTION 222**
An organization is making a strategic change to customer support in response to feedback. They plan to extend their helpline availability hours.
Why is the organization making this change?

A. Users expect professional expertise
B. Users require personalization
C. Users expect always-on services
D. Users require regional access

**Correct Answer: C**
**Section:**

**QUESTION 223**
Which technology allows organizations to run multiple computer operating systems on a single piece of physical hardware?

A. Hypervisor
B. Containers
C. Serverless computing
D. Open source

**Correct Answer: A**
**Section:**

**QUESTION 224**
An organization wants to develop an application that can be personalized to user preferences throughout the year.
Why should they build a cloud-native application instead of modernizing their existing on-premises application?

A. Developers can rely on the cloud provider for all source code

B. Developers can launch new features in an agile way
C. IT managers can migrate existing application architecture without needing updates
D. IT managers can accelerate capital expenditure planning

**Correct Answer: B**
**Section:**

**QUESTION 225**
An organization wants to digitize and share large volumes of historical text and images.
Why is a public cloud a better option than an on-premises solution?

A. In-house hardware management
B. Provides physical encryption key
C. Cost-effective at scale
D. Optimizes capital expenditure

**Correct Answer: C**
**Section:**

**QUESTION 226**
An organization is using machine learning to make predictions. One of their datasets mistakenly includes mislabeled data.
How will the prediction be impacted?

A. Increased risk of privacy leaks
B. Increased risk of inaccuracy
C. Decreased model compatibility
D. Decreased model training time

**Correct Answer: B**
**Section:**

**QUESTION 227**
An organization wants its users to validate a series of new features for their app.
Why should they use App Engine?

A. Because their app is containerized and enabled by microservices
B. Because the updated app will only include new features
C. To run different versions of the app for different users
D. To run different versions of the app for the same user

**Correct Answer: C**
**Section:**

**QUESTION 228**
An organization's web developers and operations personnel use different systems.
How will increasing communication between the teams reduce issues caused by silos?

A. By assigning blame for failures and establishing consequences
B. By combining job role responsibilities to ensure that everyone has shared access
C. By increasing data encryption to strengthen workflows
D. By emphasizing shared ownership of business outcomes

**Correct Answer: D**
**Section:**

**QUESTION 229**
An e-commerce organization is reviewing their cloud data storage.
What type of raw data can they store in a relational database without any processing?

A. Product inventory
B. Product photographs
C. Instructional videos
D. Customer chat history

**Correct Answer: A**
**Section:**

**QUESTION 230**
A global organization is developing an application to manage payments and online bank accounts in multiple regions. Each transaction must be handled consistently in their database, and they anticipate almost unlimited growth in the amount of data stored.
Which Google Cloud product should the organization choose?

A. Cloud SQL
B. Cloud Spanner
C. Cloud Storage
D. BigQuery

**Correct Answer: B**
**Section:**
**Explanation:**

**QUESTION 231**
An organization is struggling to meet user demand for change and wants to modernize their legacy applications by moving the applications to the cloud
Why would this help the organization satisfy user expectations'?

A. Toil automation helps make automatic updates
B. Updates can be pushed out more quickly to repair bugs
C. Customer data can be used to offer tailored content
D. DevOps requires that industry trends be measured and tracked

**Correct Answer: B**
**Section:**
**Explanation:**

Moving legacy applications to the cloud can help organizations satisfy user expectations by enabling them to push out updates more quickly to repair bugs.

**QUESTION 232**
An organization wants to build autoscaling web applications without having to manage application infrastructure
Which Google Cloud product should they use?

A. App Engine
B. AutoML
C. Anthos
D. Apigee

**Correct Answer: A**
**Section:**
**Explanation:**
Per Google docs, App Engine, allows for 'freeing up your developers with zero server management and zero configuration deployments'. https://cloud.google.com/appengine

**QUESTION 233**
An international bank is looking for a serverless warehouse solution that lets them perform smart analytics
Which Google Cloud product or service should the bank use?

A. BigQuery
B. Dataflow
C. Compute Engine
D. Cloud Spanner

**Correct Answer: A**
**Section:**
**Explanation:**
The international bank should use Google Cloud's BigQuery service, which is a fully managed, serverless data warehouse that allows for high-speed analysis of large datasets. It provides a range of built-in functions for analytics and can easily integrate with other Google Cloud services.

**QUESTION 234**
An organization decides to migrate their on-prenwses environment to the cloud They need to determine which resource components still need to be assigned ownership
Which two functions does a public cloud provider own? (Choose 2)
Choose 2 answers

A. Fixing application security issues
B. Infrastructure architecture
C. Hardware capacity management
D. Hardware maintenance
E. Infrastructure deployment automation

**Correct Answer: C, D**
**Section:**

**QUESTION 235**
What is an organization exclusively responsible for when they access an application through a software as a service (SaaS) model?

A. Maintaining overall system operability
B. Maintaining customer-facing content
C. Monitoring data center servers
D. Monitoring computer networks

**Correct Answer: B**
**Section:**

**QUESTION 236**
An organization is looking for a business intelligence solution that allows individual employees and end users to analyze business data and generate insights.
Which Google Cloud product or service should the organization use?

A. Looker
B. Cloud Spanner
C. BigQuery
D. Dataflow

**Correct Answer: A**
**Section:**

**QUESTION 237**
An organization cannot afford to modernize their infrastructure but they want to process data from their legacy system in a modern platform hosted by a business partner
What solution should the organization choose to make their data accessible?

A. Compute Engine
B. Anthos
C. An application programming interlace
D. Google Kubernetes Engine

**Correct Answer: C**
**Section:**

**QUESTION 238**
An organization needs frequent access to only a subset of their data. They want to reduce costs by depositing the rest of their data across Nearline Coldlme and Archive repositories
Which Google Cloud product should the organization use?

A. Filestore
B. Cloud Spanner
C. Data Catalog
D. Cloud Storage

**Correct Answer: D**
**Section:**
**Explanation:**
Per Google docs, specifically for GCP Cloud Storage there exists four types of storage with one of them, standard storage, being described as 'storage for data that is frequently accessed ('hot' data) and/or stored for only brief periods of time.' https://cloud.google.com/storage

**QUESTION 239**
A retail organization has moved all of their inventory data to a relational database in the cloud.
What functionality does a relational database offer?

A. It analyzes unstructured data which can then be accessed in multiple regions
B. It stores transactional data which can then be accessed electronically
C. It stores large amounts of raw data in its original format
D. It rapidly analyzes large and multi-dimensional datasets

**Correct Answer: B**
**Section:**
**Explanation:**
A relational database offers the functionality of storing transactional data, which can then be accessed electronically. Relational databases store structured data that can be organized in tables with defined relationships between them. This makes them well-suited for transactional data, such as inventory data, that needs to be accessed and updated frequently.

**QUESTION 240**
An organization wants to introduce a new image recognition login system
What should the organization do to follow SRE principles?

A. Roll out the new system to a subset of employees to test it out
B. Roll out the new system to all employees to collect as much data as possible
C. Avoid rolling out the new system because it may have security flaws
D. Avoid rolling out the new system because it may violate privacy policy

**Correct Answer: A**
**Section:**

**QUESTION 241**
An organization has created an application that can diagnose different medical conditions when users submit images of their affected body parts.
Which Google Cloud product or service did the organization use?

A. App Engine
B. Machine learning
C. Cloud Logging
D. Cloud Profiler

**Correct Answer: B**
**Section:**

**QUESTION 242**
An organization is struggling to keep up with the growth of their application which is running on legacy infrastructure.
What might be holding them back?

A. The inaccessibility of their data due to perimeter security
B. The overreliance on platform as a service
C. The time it takes their serverless compute function to scale
D. The cost of provisioning hardware for peak usage

**Correct Answer: D**
**Section:**
**Explanation:**
Legacy infrastructure is typically based on on-premises hardware that is managed and maintained by the organization. As the application grows and the user base expands, the hardware required to support it must also grow. This can lead to significant costs associated with provisioning and maintaining hardware, particularly if the organization needs to provision for peak usage.

**QUESTION 243**
What DevOps practice should an organization use when developing their application to help minimize disruption caused by bugs?

A. Pause production until all bugs have been eliminated
B. Prioritize fixing large bugs during production because they are easier to review
C. Implement small changes incrementally to reduce recovery time when bugs appear
D. Implement large changes together to make rolling back easier when bugs appear

**Correct Answer: C**
**Section:**
**Explanation:**
One of the key principles of DevOps is to release changes frequently and in small batches. This helps to reduce the risk of disruption caused by bugs. If a bug is introduced in a small change, it is easier to identify and fix the bug without affecting a large number of users.